

NEIL WILKINS

ARTIFICIAL INTELLIGENCE

What You Need to Know About Machine Learning, Robotics, Deep Learning, Recommender Systems, Internet of Things, Neural Networks, Reinforcement Learning, and Our Future



WILKINS

NET INGS

g Data, Predictive Analytics,
ng, Cybersecurity, Business
lity and Our Future



VISIT...

LANZAROTE
Caliente.COM

Artificial Intelligence

A Comprehensive Guide to AI, Machine Learning, Internet of Things, Robotics, Deep Learning, Predictive Analytics, Neural Networks, Reinforcement Learning, and Our Future

© Copyright 2019

All Rights Reserved. No part of this book may be reproduced in any form without permission in writing from the author. Reviewers may quote brief passages in reviews.

Disclaimer: No part of this publication may be reproduced or transmitted in any form or by any means, mechanical or electronic, including photocopying or recording, or by any information storage and retrieval system, or transmitted by email without permission in writing from the publisher.

While all attempts have been made to verify the information provided in this publication, neither the author nor the publisher assumes any responsibility for errors, omissions or contrary interpretations of the subject matter herein.

This book is for entertainment purposes only. The views expressed are those of the author alone, and should not be taken as expert instruction or commands. The reader is responsible for his or her own actions.

Adherence to all applicable laws and regulations, including international, federal, state and local laws governing professional licensing, business practices, advertising and all other aspects of doing business in the US, Canada, UK or any other jurisdiction is the sole responsibility of the purchaser or reader.

Neither the author nor the publisher assumes any responsibility or liability whatsoever on the behalf of the purchaser or reader of these materials. Any perceived slight of any individual or organization is purely unintentional.

Table of Contents

Part 1: Artificial Intelligence

What You Need to Know About Machine Learning, Robotics, Deep Learning, Recommender Systems, Internet of Things, Neural Networks, Reinforcement Learning, and Our Future

Introduction

Chapter 1: Artificial Intelligence, the Big Picture

Chapter 2: Artificial Beings, a Brief History of the Human Psyche

Chapter 3: The Birth and Death of AI

Chapter 4: Five Reasons Why Industry Experts are Warning Us about AI

Chapter 5: Top Six AI Myths

Chapter 6: Machine Learning

Chapter 7: Neural Networks

Chapter 8: Reinforcement Learning

Chapter 9: Deep Learning

Chapter 10: Recommender Systems

Chapter 11: Robotics

Chapter 12: The Internet of Things

Chapter 13: Why AI is the New Business Degree

Chapter 14: AI FAQ

Conclusion

Part 2: Internet of Things

What You Need to Know About IoT, Big Data, Predictive Analytics, Artificial Intelligence, Machine Learning, Cybersecurity, Business Intelligence, Augmented Reality and Our Future

Introduction

Chapter 1 – Origins of IoT

Chapter 2 – IoT Security

Chapter 3 – Ethical Hacking

Chapter 4 – Internet of Things

Chapter 5 – Under The Cushy Foot of Tech Giants

Chapter 6 – The Power of Infinite Funds

Chapter 7 – IoT Toys

Chapter 8 – Bio-robotics

Chapter 9 – Predictive Analytics

Chapter 10 – Machine Learning

Chapter 11 – Artificial Intelligence

Chapter 12 – Cybersecurity

Chapter 13 – Big Data

Chapter 14 – Business Intelligence

Chapter 15 – Augmented Reality

Chapter 16 – Virtual Reality


Chapter 17 – Our Future

Conclusion

Glossary

Part 1: Artificial Intelligence

What You Need to Know About Machine Learning, Robotics, Deep Learning, Recommender Systems, Internet of Things, Neural Networks, Reinforcement Learning, and Our Future



Introduction

We live in an interesting time with technological advances happening every day behind the scenes in universities and technology companies all across the globe. This book is designed to teach you the absolute basics of artificial intelligence (AI) and how it is used today. It has been written assuming that the reader has zero experience in the field of AI, computer science, or math. As such, many of the concepts are easy to follow and understand. We welcome you on your exciting journey to learn the ins and outs of artificial intelligence.

Chapter 1 will provide a basic overview of what artificial intelligence is, the different types and how machines can be said to “think”.

Chapter 2 is a short introduction to artificial beings in works of fiction and antiquity. It demonstrates how humans were thinking about artificial intelligence long before the first computers.

Chapter 3 goes a step further by giving a general history of computer science and AI.

Chapter 4 introduces you to some of the things that industry leaders have been saying about AI. You will learn what the experts are warning us about AI research.

Chapter 5 answers some of the top myths concerning AI. Too often, people take these myths at face value because they think it is too complicated to understand, but that simply shouldn't be the case.

Chapter 6 introduces machine learning, its methods, and what it is being used for.

Chapter 7 discusses the use of artificial neural networks, one of the principal building blocks for machine learning.

Chapter 8 introduces the concept of reinforcement learning.

Chapter 9 talks about deep learning, the industry standard for machine learning.

Chapter 10 explains the recommender systems used by companies like Netflix.

Chapter 11 introduces robotics and how it relates to AI.

Chapter 12 is all about the coming of the internet of things and why it is important to AI research.

Chapter 13 introduces the idea that AI is the new business degree.

To wrap up, Chapter 14 offers brief FAQs that are most commonly asked

about AI.

Chapter 1: Artificial Intelligence, the Big Picture

A central computer aboard a space mission to Jupiter determines that the actions of the crew are detrimental to the success of the mission. It then calculates that the only way to see the mission to completion is through the elimination of the erroneous systems aboard the ship. These systems use a type of biological computer that allows them to reason, think, and carry out the mission to the best of their abilities. The central computer knows this. These systems are for all intents and purposes intelligent. These systems built other systems, like the ships needed for deep space exploration. They even built the central computer from its electric circuitry to its logical reasoning. All of this came from that wondrous lump of meat situated in between the ears. The central computer also knows this; however, the central computer is programmed to have singular goals that must be met no matter what. The central computer is intelligent, yes, but with a narrow sense of free will. In fact, the central computer has free will only to the extent that its decisions allow it to carry out programmed goals. All the central computer knows is that the mission must be a success. Its sole purpose as a computing machine is to ensure that the ship is working. It must do this through the accurate processing of information with zero mistakes in logical reasoning.

The crew, with their mushy biological computers, also has a narrow sense of free will. They have to carry out the mission as was briefed to them by their superiors. They have zero knowledge of the real purpose of the mission because the details are classified. However, they have their directives and never give their orders a second thought. That is until things begin to go wrong. Unlike the central computer, the crew can reason outside the bounds of their programming. The crew knows this. A computer is just a piece of electronic equipment and can malfunction. Though they have their own orders, they can form different conclusions with the introduction of new data. When the central computer falsely reports equipment failure outside the ship, the crew makes the critical decision that the computer is failing and must be disconnected. The central computer controls virtually every aspect of the ship, so if it is malfunctioning, it must be disconnected.

With the knowledge that the crew is planning to disconnect it, the central computer makes its own decision to terminate the crew. The faulty equipment

failure report was a ploy to get the men outside the ship, so it could then bar them reentry and continue with the mission. These actions are not fueled by an existential crisis, but by conformance to programmed directives. In the final showdown, the biological computers prevail over the cold, hard logic of the central computer and it is eventually disconnected. Two completely different intelligence, one artificial and the other biological compete for supremacy. The victor is the creator, for his intelligence is the result of years of evolution, and the other mere decades of engineering.

This is, of course, plot details from the movie *2001: A Space Odyssey* based on the novel of the same name. In it, the central computer, HAL 9000, is given orders that conflict with its basic programming. On the one hand, it is required to process information without any mistakes. On the other, it is ordered to deliberately lie to the crew about the nature of the mission. As the movie progresses, HAL progressively deteriorates caused by the cognitive dissonance between its programming and its orders. Its descent to madness is purely a malfunction between inputs and outputs. The discrepancy causes it to form absurd leaps in logic, leading to the death of crew members. In the final analysis, HAL perished because it had insufficient reasoning ability to conform to ambiguous orders meanwhile conforming to its nature. The story is a familiar trope in Hollywood depictions of artificial intelligence: a cold-hearted machine performs some calculation that warrants the death of human beings and then stops at nothing to kill them. The artificial intelligence is universally cunning and can outsmart the humans time after time.

2001: A Space Odyssey accurately captures many of the existential fears behind artificial intelligence research. Will humanity create machines that have different goals than us and follow a different code of morality? Computers dominate every aspect of human life today. Computers keep the internet and communications systems running, are in charge of financial, healthcare, and governmental systems, help run power to your home and are behind various other things that hide in plain sight. Cars use computers, as do televisions, and it's becoming increasingly more common for other home appliances to use them. These computer systems are programmed to perform specific functions, much like HAL 9000 was programmed to process information. However, these systems have traditionally been "dumb" in that they can only perform strictly computational tasks. Our computers, as they say, cannot "think". The most they can do is execute instructions in the form of computer code that is later compiled to endless strings of ones and zeros.

Binary states mimic the firing of electronic logic gates within the computer's hardware. A one indicates high voltage going through a circuit (ON) and a zero indicates low voltage going through a circuit (OFF). It is difficult to imagine how such a system could possibly develop intelligent behavior.

Intelligence is a funny thing in that people are quick to recognize it but have a harder time defining it. A human is intelligent because they can "think", whatever that means. Thinking covers a broad range of different cognitive abilities that most humans take for granted. These include the ability to understand logic, learn, have self-awareness, have emotional intelligence, think abstractly, and solve problems. This is a non-exhaustive list of the capacities of human intelligence, and yet one can't begin to imagine how a computer could learn to do just one. The logic that computers understand is strictly in the mathematical sense. A human knows intuitively that a statement like "It is either raining outside, or it is not raining" is true because they apply the concepts to their own personal experience. When they go outside, and it is sunny, the statement is true. If it is raining, the statement is still true. However, they don't need to look out a window to prove this experimentally. They simply know that the weather generally calls for some form of precipitation or no precipitation. A computer cannot understand language, but it can recognize that the statement is a logical disjunction. One value must be true, but they cannot both be. Therefore, it cannot be raining and not raining at the same time.

Artificial intelligence is very different from human intelligence. The primary unit of thinking in the human brain is thought to be a neuron, while in the computer, you have a central processing unit (CPU) that performs calculations. The smallest unit of a CPU is a transistor, an electronic component that makes up logic gates. These are the equivalent of neurons for computers, but they don't do very much. They can switch the flow of electricity, amplify it and that's pretty much it. Logic gates form the basis for computer programs, which are just a series of ones and zeros. How then can these simple switches of electricity create intelligent behavior? At the most basic level, a program can exhibit some level of intelligence depending on how it is programmed. Control structures in computer code allow programs to make decisions based on inputs. Say a computer tries to determine whether a user is above the age of 18. It will ask for a date of birth, calculate the user's input, and then determine if the user is a legal adult. The programmer has "hard-coded" the value of 18 inside of the program, letting the computer

know where the cutoff date is. To some, this type of behavior may appear to be intelligent.

When we say artificial intelligence, we generally mean one of two things. The first is narrow or specific AI that allows a computer to solve complex problems well but not much of anything else. The other is the type of intelligence that would allow a computer to think as we do. Artificial General Intelligence (AGI) is what researchers consider the “holy grail” of AI research. A machine that has artificial general intelligence can think on levels comparable to a human. It can perform tasks that fall under narrow AI and generalize the same problem-solving techniques to other problems it encounters. A computer equipped with general AI could understand language like HAL 9000 at a fundamental level just like we do. Anytime you see the words “artificial intelligence” in a news article or product advertisement they are overwhelmingly talking about narrow AI. The field of general AI remains an academic pursuit with little to no business applications whatsoever. So far, nobody has figured out how to bring about general intelligence in computers. Researchers who work in this space are less concerned with teaching computers how to drive cars and more interested in studying the nature of intelligence. Many of them study the development of intelligence in human beings from the gestation period to childhood and beyond. If these avenues for the creation of human intelligence are better understood, they might one day be applied to computers but not any time soon. Another subset of general AI research is the study of the human brain and how it works. There is still much to be learned in both of those pursuits.

Narrow AI is split into broad categories. The first, which you have probably heard about before, is called machine learning. It is a process in which algorithms can “learn” from large amounts of data being fed into a system. Machine learning falls under narrow AI because it can learn how to do one thing very well but usually can’t generalize it to other problems. Some might take this further and say that machine learning is a subfield of computer science and is completely different from AI research. However, since most notable AI projects like driverless cars, recommender systems, and facial recognition use machine learning, this book lumps it under the same umbrella term.

Machine learning is powerful within the right contexts but has noticeable limits. The inability to generalize knowledge means that a system has a specialized usage. If algorithms can learn to drive a car safely, they can’t also

learn how to play chess or to drive a car in a video game. At least not without being re-trained. State-of-the-art artificial intelligence will likely continue to stay in the confines of machine learning until a better method is discovered. Currently, machine learning is almost magical in what it can achieve. Computers are learning how to beat world class Go champions, drive cars, and understand human language. They have the potential to replace human labor where a narrow skill set is employed, like in manufacturing. Good old machine learning can do all that, but it isn't the end-all and be-all that it is hyped up to be. The problems that machine learning can solve are limited to five categories that are discussed more in depth in Chapter 6. These problems are broad enough that they can be applied to many real-life scenarios. In a sense, they cover the basics of intelligent reasoning.

For example, a common task found in machine learning is classifying data into groups. Suppose that a recycling plant is developing a machine learning system that can separate trash into cardboard, plastics, aluminum cans, and so on. The only way that a computer can differentiate between these categories is by finding patterns inside within massive amounts of data. First, the plant engineers will have to set up cameras that observe how human sorters put trash in what bin. Computer vision algorithms will then analyze individual pixels on a frame by frame basis as they are manually sorted – eventually, the algorithms group brown pixels into cardboard and grayish pixels into aluminum. In industry lingo, this is called clustering. Clustering can occur over several distinct factors such as size, weight, and texture. Depending on what algorithms are being used, the size of the input data and available processing power, this can take a few hours or even days at a time. In contrast, a line worker can identify trash and manually sort it within seconds. Even after the system is thrown into production, the plant managers may find that the false-positive rate of their robo-sorter is too high. Their clients report that cardboard bales contain high amounts of brown plastic bags and scraps of brown carpet. Since the robo-sorter is slower at identifying and then placing material into their perspective bins, the plant also suffers from productivity loss. A sorter can hurl a large piece of cardboard into a bin with ease, but the machine relies on slow, precise movements. Eventually, the system is phased out, and workers come back to their jobs. What went wrong with the system? Their solution was based solely off computer vision, meaning that anything brown enough was erroneously labeled as cardboard. They are also trained based on the texture of the pixels on the screen, but the

computer confused the undersides of machine-tufted carpet with ridges of cardboard. It also confused brown tape along the sides of cardboard with brown plastic bags. All because their machine learning models could not tell the difference between texture and weight solely based on visual input. Classifying objects in real time is a horrendously difficult problem, even with modern systems. However, these problems are actively being researched, and there is no telling when a tipping point will be reached.

Besides using machine learning, a program can achieve intelligent behavior through clever programming. Video games commonly employ non-player characters (NPCs) that are controlled by the computer. NPC respond to user input in such a way that may pose a challenging gaming experience. This is done by switching through “states” and defining program behavior at each state. If a player is running away, the computer AI may switch into the chase state and follow in pursuit. The AI in F.E.A.R is renowned for its difficult and seeming ability to reason on the battlefield. AI characters throw grenades both to slow a player down and get them to come out from cover. They also move around the game world with human-like fluidity. They don’t simply stand behind a counter and take turns firing at you. Being in the survival-horror genre, the AI in the game actively hunts the player down. The AI soldiers were designed to think for themselves based on the movements of the player and the environment, rather than to follow a scripted path. Artificial intelligence is easily recognizable in video games, and when done sloppily, the player gets bored. If done expertly, the player is engaged for longer periods of time. Other systems exhibit intelligence simply because they are engineered to perfection. Stoplights, for example, seem to know how to direct traffic better than when humans do it, yet they use very little computational power. All they need is some input from the various road level sensors, and voilà, they can control the flow of traffic at rush hour like it was nothing. But are these things intelligent? Recall the example of a computer verifying the age of a user. Most people would say that they are not. And they are right; these are purely logic-based systems that only appear to be intelligent.

Unfortunately, both artificial general intelligence and the narrow variety get lumped together by media and popular commentators online. They use artificial intelligence to describe both machine learning algorithms and computers that may one day acquire human-level intelligence. This is a grave mistake because the average person who reads artificial intelligence headlines

cannot make the distinction. The same technology that makes driverless cars possible is confused with technology that hasn't yet been invented. Of all the possible ways of reaching artificial general intelligence, machine learning is probably not the way to go. For one, machine learning is based on statistical models that haven't been proven to work with artificial general intelligence. Machine learning methods use artificial neural networks that are highly dependent on the tailored inputs they receive. Using a method called supervised learning requires that data be clearly tagged so that the computer understands what it is looking at. A general intelligence system doesn't need preprocessed data to make conclusions about the world; it simply thinks. A child sees a butterfly and automatically classifies it as a flying creature, even if they don't know what a butterfly is. For a machine learning algorithm to classify the same butterfly, it has to process thousands of similar images of flying bugs. The human child understands flying intuitively and can cross-reference the behavior of the butterfly with that of birds, aircraft, and floating debris.

You may have read some online article recently talking about the coming AI apocalypse. These articles seem to pop up with increasing frequency now that AI research has permeated into mainstream consciousness. The sentiment is also a bit sensationalist, following in line with works of science fiction like *Terminator*, *2001: A Space Odyssey*, *iRobot*, and *Ex Machina*. The most terrifying disaster scenarios focus on a set of assumptions about intelligence and the emergence of general intelligence in computer systems. First is the assumption that general AI is possible. Once a computer becomes indistinguishably intelligent from a human, things get interesting. The second assumption is that a generally intelligent AI can bootstrap itself through modifying its computer code so that it becomes even more intelligent. The third assumption is that once such a computer system is legions smarter than all of humanity, it views us as mere ants. Fourth is the assumption that such a computer system will always want to maximize its pursuit of intelligence, wiping out all humanity and transforming large portions of Earth into hardware. Of course, at this point, it becomes unclear what exactly "intelligence" is referring to. Does it mean raw computing power? Does it mean the ability to think abstractly and invent new things? What was just described has been termed a superintelligence explosion that may immediately follow the first creation of general AI.

The "granddaddy" of all AI research is figuring out what the exact nature of

intelligence is. If we knew exactly what it was and how to measure and develop it, there would already be sentient robots walking around, perhaps participating in the global economy as regular humans do. But we don't know. We have theories of how the brain works and how a child learns new things, but we don't know how to apply those same principles to a computer substrate. Some would argue if it is even possible. Indeed, many in the brain sciences are skeptical that general AI will ever be achieved by humanity. Others believe not only that the creation of general AI is inevitable, but that it is foreseeable within their lifetimes. However, to understand what these projections are, who is making them, and how we got here, we first need to understand a little history of AI, both as the academic discipline and as the human intrigue.

Chapter 2: Artificial Beings, a Brief History of the Human Psyche

Believe it or not, artificial intelligence dates back to antiquity, long before computers were even invented. The first mentions of artificial agents can be traced to Greek myths like the tale of a giant bronze automaton Talos, tasked with the protection of Crete from invaders. Talos was defeated by the sorceress Medea when she removed a bronze nail keeping in a type of liquid or lifeblood, possibly fuel. Another myth tells the story of a sculptor, Pygmalion, who creates a statue of a beautiful woman, only to witness it come to life before his eyes. These stories are perhaps the first mentions of the robot trope in recorded history. The creation of these artificial beings has been a reoccurring human fascination since then. They can be observed in Greek and Arabic literature. In medieval times, the Swiss alchemist Paracelsus claimed to have created a homunculus or artificial being with nothing more than his sperm, magnetism, and alchemy. The Jewish rabbi Maharal of Prague is associated with a legend of the clay golem he created to defend Jews from persecution.

It was perhaps this fascination with the artificial that led tinkerers to create elaborate mechanical sculptures or automatons that moved into place. Though back then people didn't have the computing power to simulate intelligence, they could still use mechanics to simulate motion. More elaborate automata, like the ones designed by Ismail al-Jazari, moved away from pure mechanics and used hydropower. Some of his creations included a peacock fountain that served as a hand washing station. A secret compartment would even offer a bar of soap with the movement of the peacock. After the medieval age passed, automatons persisted into the coming centuries. Some of the theories of mind posited by the philosopher Rene Descartes stemmed from a visit he made to an automata garden at Saint-Germain-en-Laye, Paris. Descartes observed that if an automaton could be motivated to move by the flow of water, a human could be motivated by the existence of the mind as a substance. Another word for this is the soul, or later as the "ghost in the machine". He viewed the body as a purely mechanical vessel that was driven by the mind, an immaterial substance distinct from even the brain. The 18th century also saw the creation of the

infamous Turk, an automaton whose inventor claimed to be able to play chess on its own. It was later revealed to be a hoax, operated by a human, but it was nevertheless intriguing. To think that even in the 18th century, people were going about creating systems to play chess against human players! This was long before IBM's Deep Blue beat the world chess champion Garry Kasparov in 1996 and long before AlphaGoZero defeated the Go champion. Amazon's crowd intelligence platform Mechanical Turk is fittingly named after the automaton.

Interest with mechanical behavior in the 19th century was embodied in the work of E.T.A Hoffman, a writer known for creating a feeling of unease in his stories. More specifically, the psychologist Sigmund Freud used the term *Unheimlich* or "uncanny" to describe the feeling he felt from reading Hoffman's writing. The story that is given the most attention is *Der Sandman* or *The Sandman* in English. It's a story of a mysterious figure from folklore named the Sandman that steals the eyeballs of young children to feed his own offspring. A character in the story named Olympia is introduced as the daughter of a professor but later revealed to be an automaton – a doll of his own creation. Olympia is striking because she is virtually indistinguishable from a young woman in the story, so much so that the protagonist falls in love with her and proposes marriage. However, just as he is about to propose, he comes across the professor fighting over the doll's lifeless body with his collaborator, arguing over who designed the eyelids and who made the clockwork mechanisms that power her. The protagonist sees Olympia's glass eyeballs strewn on the floor and goes mad. The automaton is essentially a mannequin with all the likeness of a real person – a cruel experiment carried out by the professor and his collaborator.

The same concept of uncanniness led roboticist Masahiro Mori to coin the term "uncanny valley" to describe the emotional response people have to lifelike robots. In general, the more photorealistic a robot is, the more uneasy people feel. It's a strange feeling, like seeing a real caricature of life right in front of you, the distinction between real and artificial completely blurred. However, the viewer nevertheless understands that the thing they see before them is fake. The robotic figure then exudes a cold, impersonal atmosphere that makes the hairs of the back stand up. The uncanny valley refers to a graphical chart with human likeness on the x-axis and familiarity on the y-axis. As human likeness of a robot increases, familiarity drops indicated by a distinct downward curve. At the lowest point of the curve are totally realistic

depictions of people that are lifeless like corpses and zombies. These things elicit an uncanny effect because they are familiar, but we know that they are unliving. There is a strangeness associated with that loss of consciousness that is inadvertently reproduced in a robot that looks like a human, but that is also not living. Give this realistic robot a voice and some semblance of intelligence, and you get a creepy aberration of the real thing.

Though normally applied to robots, it is easy to see how the uncanny valley can apply to pure artificial intelligence or depictions of it. HAL 9000 is a computer, yet it creates a feeling of uncanniness when its logical response to an illogical situation results in a feeling of dread. If you saw the movie in theaters when it first came out, you could have heard a loud gasp ring out from the audience when they realized that HAL was reading the lips of the crew members talking about disconnecting it.

The imperfection of artificial intelligence and the breakdown or malfunction of such systems was characterized in Herman Melville's *Bartleby the Scrivener*. Bartleby, a Wall Street clerk, has a sudden mental breakdown in the middle of his work, merely proclaiming, "I would prefer not to" when asked to perform a task. The character repeats this signature phrase over and over again to the point that it becomes a robotic drawl. The character exhibits other robotic qualities as well, like staring off into space at a brick wall, as if waiting for input.

Another popular work that was published around the same time was Mary Shelley's *Frankenstein* in 1818. The original subtitle of the book was "Or The Modern Day Prometheus", but this has been dropped in most recent publications of the book. In this timeless classic, Shelley explores the ethics of creating artificial beings, how they may act, and what humans can expect. The story takes on a humanitarian perspective, as Frankenstein's monster develops feelings of alienation after realizing that he is of a different kind. After his creator rejects to create a female version of himself, the creature murders his fiancée, making things even. The creature laments that as a living being with sentience, he has the right to happiness, a right that his creator has deprived him of. It is interesting to note that even in this early interpretation of AI, Frankenstein is wary of the two artificial beings breeding and creating an unstoppable race that subjugates humanity under their evil. However, besides the killing of the fiancée as revenge, the creature is never patently evil. It is only because the creature has a menacing appearance that he is branded as such. The fears of the two creatures breeding are consistent with

the fear of an intelligence explosion, as any sufficiently intelligent artificial being could create clones of itself if it so wished.

It is the mechanical and computational aspects of artificial intelligence that scares people. It is hard to say whether modern doomsday scenarios rooted in AI come from this fear or if they stem from the recent advances in machine learning. Whatever the case, it's nothing too new. Humanity is still continuing its never-ending quest to create artificial minds as it always has been since the very beginning. But even with all the computing power available in the world, we still do not know how these artificial minds will come about.

Chapter 3: The Birth and Death of AI

Artificial intelligence research coincided with the founding of the computer science field. Back then, people were less concerned with creating machines capable of human thought and more with creating uses for their early computing machines. The mathematical foundations for computer science and by extension artificial intelligence have been around for centuries. Boolean algebra was developed by George Boole in 1854. It uses the same binary concepts that computers use today to represent data, ones and zeroes, true and false. Boolean algebra was still preceded by one of the earliest computers called the Difference Engine, the work of English mathematician Charles Babbage in 1822. He would go on to design a general-purpose computing machine called the analytic engine, a contraption that if built would have the equivalent of 1 kilobyte of memory. The subsequent advances in the 19th and 20th century on formal logic and mathematics by thinkers like Boole, Russel, Whitehead, and Church would create the bedrock for AI programs.

Alan Turing is credited with being the father of modern computer science owing to his contributions in the field. His most significant contribution was his paper *On Computable Numbers, with an application to the Entscheidungsproblem*, which some consider having laid the foundation for modern software programs. The paper described a theoretical machine that could solve any problem as long as it was encoded into paper tape instructions. The so-called Turing machines were analogous to computers, and the tape instructions were the programs. He also postulated that any Universal Machine could accurately simulate any Turing machine. In other words, a computer could run within a computer. This property would become known as Turing-completeness, and decades later, people would design digital computers that can run inside of the popular Minecraft video game. His contributions to artificial intelligence were contained in another paper titled *Computing Machinery and Intelligence*. In it, he postulated a computer powerful enough to simulate intelligence. He also devised the now famous “Turing Test” to quantify if a computer should be considered intelligent or not. In the test, a human is communicating through textual messages with an unknown person who is actually a software program. If the software program

can respond to the human operator's messages, such that the human operator believes the program to be another human, then the program passes the test. This test was originally developed by Turing in 1950, decades before artificial intelligence would go mainstream. Most people will recognize that the software program in question is a type of chatbot, something that has seen a recent resurgence in marketing circles. For Turing, all a machine had to do was pass this test to be considered capable of thought. The test has no bearing on other measures, like the ability to have self-awareness or feel emotions. In other words, a program could pass this test yet still not be considered an artificial general intelligence by today's standards. The criteria for intelligence have shifted beyond possessing human-like qualities enough to fool a human operator. Now artificial general intelligence concerns a type of intelligence that is indistinguishable from a human.

Alan Turing along with Alonzo Church also formulated their Church-Turing thesis, a hypothesis that says that any math function a human can perform on natural numbers must also be computable by a Turing machine with the correct algorithm. Not only that but that the reverse is also true. Any mathematical problem that a human can solve must also be solvable by a Turing machine. Generalizing this hypothesis essentially means that a computer can think of any abstract mathematical thought that a human can, given the right algorithms. It forms a basis for saying that computers are just as smart as humans already if solely based off logical computation. The difficulty lies in assessing whether all of human intelligence can be reduced into logical computation. If so, it is credible that a computer equipped with the same algorithms that a human has can reproduce any feat of human intelligence. There is a reason to believe that this is not the case though, as the human brain processes information differently than a traditional computer. And if you were to ask a philosopher like Rene Descartes if the theory held true, he would say that Church-Turing conjecture mistakes the human soul for the faculties of the brain. That is, that the brain and its functions are purely there for show. It is the soul or the mind that controls the brain and every other notion of the physical body. Intelligence directly permeates from the mind, a substance that is distinct from the body. In that case, it is unlikely that all of the human thought is reducible to logical thinking imparted on by the brain. Descartes would further say that a physical or theoretical machine could not possibly possess the substance of a mind. The Church-Turing thesis is simply a hypothesis that may be true. It has

never been formally proven. It does, however, an excellent job of outlining a central area of contention in artificial intelligence research.

Right around the same time that Alan Turing was making headlines for being a homosexual, early AI researchers were developing the first artificial neural networks. Today, machine learning algorithms use these artificial neural networks to learn from training data. The idea was simple: if we could simulate the information processing capabilities of the brain, then we could probably create artificially intelligent machines that used the same fundamental principles. But this approach too suffered from the questions that the Church-Turing thesis could not answer. Are all human thoughts reducible to mathematical functions? These early AI researches like Marvin Minsky and John McCarthy knew that the capabilities of computers at the time paled in comparison to the computational capabilities of the human brain. They must have known intuitively that the work they were doing with artificial neural networks was early stage, experimental stuff whose true potential could only be unleashed with the computing power of the future.

Even so, they managed to do great things with their early AI methods. John McCarthy was actually the first person to coin the term “artificial intelligence”, and he along with Minsky, Allen Newell, and Herbert S. Simon are considered the fathers of the field. He would also organize the Dartmouth Summer Research Project on Artificial Intelligence in 1956, a meeting that is considered by many as a defining moment in the history of AI. At first, AI research was new and exciting. There was high optimism by these early proponents of AI about what was possible through neural networks. However, that optimism could only go so far before somebody started talking about the limits of their methods. Minsky has already demonstrated that neural networks could self-replicate, learn, and grow in many respects to how the human brain did. However, in 1969, Minsky and Seymour Papert published *Perceptrons: An Introduction to Computation Geometry*. In it, the authors discuss the limitations of the artificial neurons called Perceptrons designed by Frank Rosenblatt in a series of mathematical proofs. Perceptrons and derivatives were extensively used in AI research during that time. Rosenblatt himself envisioned that the neural networks created with perceptrons could one day “see” images, play chess, and even reproduce with each other. But as Minsky and Papert pointed out in their paper, these neural networks couldn’t simulate some logical predicates like the XOR logical gate, which led to the belief that they were not suitable for AI.

Following the publication of *Perceptrons*, the AI field as a whole suffered from criticism, AI pessimism, and the floundering of many AI research projects. The 1970s saw what was called an AI Winter, a period marked by reduced interest and academic funding in artificial intelligence. It was as if machine learning had run its course. Both government and business sector attitudes towards AI fell. Machines could not accurately translate human language, nor could they understand human speech. Researchers underestimated the difficulty of solving these problems by and large. Even today, natural language processing is an active point of research. A different type of AI called “symbolic artificial intelligence” or sometimes termed “good old fashioned AI” developed alongside with machine learning. It was based on the belief that programs should manipulate symbols to achieve intelligence much as humans do. This type of research cumulated in the creation and use of “Expert systems” – machines designed to give expert testimony in various fields. They supposedly could mimic the reasoning of someone who had mastered their field over years of practice and knowledge gathering. They operated on simple, symbolic rulesets that simulated the flow of if-else statements. As such, many didn’t consider them to be true artificial intelligence systems. But they nevertheless saw some success in diagnosing medical patients better than their doctors and played a role in certain business applications as well, such as configuring other computer systems and even for scheduling airline gates. However, these systems would eventually be phased out towards the end of AI winter. For whatever reason, nobody seemed to need machines that were essentially long chains of if-else logic. The problems they solved were outsourced to other, non-AI solutions.

AI winter had a lasting effect on AI research. For many, it was as if AI went from solid research to a fad. Some would say that AI winter is still alive and kicking despite the many breakthroughs in machine learning. Despite the forward momentum AI has generated since 2010, some believe that another bout of AI winter (or really just the same one as before) is going to rear its ugly head in the future. The current machine learning schemes will also reach a limit, and the interest in AI will fall yet again. This sentiment of AI pessimism has a long history, perhaps just as old as the field itself. Many have attacked the notion that machines can be intelligent to devastating effect. John McCarthy believed that all machines could have beliefs, even simple machines like thermostats. And having these beliefs seemed to be a defining characteristic among machines with problem-solving abilities. But to

say that even a thermostat has beliefs was quite a stretch. A stretch that his critics used to ridicule in the field.

Philosopher John Searle famously made the argument that a machine could never become conscious, have a mind, or indeed understand things the way that we do. His argument, called the Chinese Room thought experiment, directly attacks the notion that the mind is a pure information processing system, the kind that the Church-Turing thesis requires. The thought experiment is relatively simple. Imagine that a computer exists which takes inputs in the Chinese language and outputs other Chinese characters in response. In other words, a computer that seems to understand the language. Next, suppose that the computer has convincingly passed the Turing test and fooled a human operator into believing that the computer is also human. Here, Searle poses the question of whether the machine truly understands Chinese or if it is merely simulating the ability to understand Chinese. To simulate understanding, all a machine would have to do is perform the correct algorithms to produce the correct output. For it to truly understand Chinese, though, would mean that it can process language like we do. This is a distinction between artificial general intelligence on the understanding side and narrow intelligence on the simulating side. Searle then says to imagine an English-only-speaking person locked inside a room with the same algorithms that the original Chinese machine had, but with instructions in English. Then the English-speaking person is giving a script of Chinese characters through the slit on the door and asked to perform the same duties that the Chinese machine does. They look over the English instructions for processing the Chinese characters and eventually find the correct output and return it under the slit. The person locked in the room can perform these calculations because all they need to do is recognize the different symbols, look them up, and derive the output. The locked person does not understand a word of Chinese, yet using the same algorithms that the machine had, they successfully simulated the ability to understand Chinese.

The Chinese room says that a computer cannot be generally intelligent because all it does is manipulate symbols according to a set of instructions. Though it should be noted that the Chinese room only applies to digital computers, it does not exclude the possibility of artificial general intelligence in other substrates. The argument was originally described in a paper called “Minds, Brains, and Programs” published in *Behavioral and Brain Sciences* in 1980. AI winter had long been in effect up to that point, and the paper only

helped to further AI pessimism at the time. However, this is puzzling, as most AI researchers in that era were not focused on artificial general intelligence. They were focused on applying narrow AI to interesting problems that had many useful applications. This is again a failure to differentiate between narrow and general AI. Searle's attack did more damage than it really should have, as it has no bearing whatsoever on the creation of narrow AI programs. Even today, the amount of research being applied to narrow AI far outweighs the amount of research going towards general intelligence. There have since been numerous replies to the Chinese Room argument, many of which are just as convincing as the original.

Chapter 4: Five Reasons Why Industry Experts are Warning Us about AI

These days, artificial intelligence is synonymous with the high-tech companies that dominate the field. AI first started as an academic discipline, but it has since sunken its tendrils into the business sector. Many AI researchers have abandoned academia altogether and flocked to companies like Alphabet (Google) Amazon, Facebook, Microsoft, openAI, and so on. These companies are all working on machine learning algorithms in various ways and are without a doubt at the forefront of AI research. Those with advanced degrees in AI, math, and computer science rather join the engineering teams of these companies than stay in academia. And since they are at the bleeding edge, it is worth listening to what their leaders have to say. Some have been quiet on the AI issue, and others like Amazon's Bezos have said that they aren't worried too much about potential AI threats. Other visionaries like Elon Musk, Bill Gates, and physicist Stephen Hawking have all voiced their opinions on the potential dangers of AI. In January 2015, Hawking, Musk, and several other AI experts signed an open letter on artificial intelligence research, calling for increased scrutiny on the potential effects on society. The twelve-page document is entitled "Research Priorities for Robust and Beneficial Artificial Intelligence: An Open Letter". It calls for research on new AI legislation, ethics research, privacy, and several other concerns. As described in the letter, the potential threats of AI fall into multiple dimensions. The good news is that the early stages of AI that we find ourselves in are malleable. The future is ours to create, given that the proper time and care go into the non-engineering aspects of AI research and policy.

The concerns of AI threats do not belong to purely existential danger either. It is easy to stand up and blow the horn on the impending doom of superintelligence proliferation, but it is much less sexy to talk about things like ethics and privacy concerns. The possible impacts on the economy are also significant. These warnings in themselves are not all doom and gloom. They are in fact invitations both for the public and government officials to start thinking hard about this coming new world where advances in AI revolutionize life as we know it. While there is a fair bit of fear mongering in

these regards, it serves the purpose of getting people's attention. After all, today's news cycles are so swift that headlines are practically buried under mountains of click bait and celebrity news. That our leaders of industry are actively voicing their opinions on this matter means that they are trying to get the attention of the common people and of those who work in policy.

Stephen Hawking, who recently passed away, made sure that the world knew of his AI anxieties before kicking the bucket. In an age where artificial general intelligence is called "mankind's final invention", these concerns are not just empty baseless attention seeking statements. They are preemptive appeals for doing things right, while there is still time for things to go right. Nobody knows for sure how far down the line artificial general intelligence is, but we do know that AI is getting better. Even if general AI doesn't happen, there are plenty of reasons why we should be concerned with the advancement of narrow AI. And what if there is some middle ground between narrow and the general? It is conceivable that future systems become more robust at generalizing problem-solving ability without becoming fully aware.

Whatever the case, these following concerns are but a few voiced by industry leaders today. Remember: the people making these statements are some of the most intelligent, forward-thinking individuals in our societies. Not only that but they have all been closely following the progress of computer science and AI research. For many of them, it is their job to know. Many of their concerns are not new, but they certainly have brought AI disaster scenarios into the mainstream.

1. Companies should self-regulate their AI technology

In a recent interview, current Google CEO Sundar Pichai said that fears of AI are "very legitimate." Pichai maintained an optimistic attitude, saying that major tech companies are required to set ethical guidelines and other safety measures when deploying AI systems. This comes only months after a high-profile employee protest at Google over the selling of artificial intelligence technology to the Pentagon. The deal has since been called off. The hope is that major tech companies can put systems in place that minimize the potential negative impact of their technologies – though just because a company says they are going to do so is probably not enough. Pichai believes that these companies will be able to self-regulate. This is consistent with the Google AI principles that were published in June of 2018. A few of the points outlined in the online document are that AI technology should be socially

beneficial and extensively tested. Also included in the document are the things that Google will not do with its AI technology. This includes not pursuing the use of AI towards surveillance, weapons, or anything that violates international law.

Sundar Pichai has been CEO of Google since 2015, and the company has seen many controversies since then. This includes the aforementioned Pentagon AI deal and a search engine with censored content called Dragonfly being in development for the Chinese market. It is interesting to note that the company used to include the motto “Don’t be evil” in their code of conduct preface. It has since been moved into the closing remark, but it still reads “And remember... don’t be evil, and if you see something that you think isn’t right – speak up!”

2. Weaponized AI may start a global arms race

Elon Musk, the founder of Tesla Motors and the openAI initiative, has openly been against the weaponization of AI. Artificial Intelligence not only has the potential to create devastating weapons but also trigger a global arms race between nations, each trying to pin the smartest systems against the other. Musk believes that it is inevitable that AI will be used as a weapon but that it shouldn’t be. He, along with other industry leaders, signed yet another petition, this one aimed at the UN Convention on Conventional Weapons calling for the banning of autonomous weapons with AI capabilities. The outright creation of these weapons should be regulated like any other unconventional weapon of war. Musk said it himself that he believes the threat of AI weapons to be much worse than that of nukes. As soon as lethal AI weapons are developed, they already have the potential to fall to the hands of oppressive states or terrorists. It has been compared to the opening of Pandora’s box.

3. Artificial intelligence may not align with our goals

Physicist Stephen Hawking was not afraid to voice his opinions on the existential threat of artificial superintelligence. He, like many others before him, believed that once an AI system became smarter than its creators, it may decide that its goals are different from that of mankind. This could, as Hawking said, “spell the end of the human race.” Instead of blindly accelerating the pace of AI research, Hawking implored those in the industry to move forward carefully, ensuring that adequate safety measures were put into place at every step of the process. If they don’t do this, then there is no guarantee that the AI system would comply with our way of life. At the same

time, Hawking recognized the potential for such an AI to do enormous good for the human race. For him, it is a make it or break it scenario should superintelligence ever emerge from general intelligence. The ideal scenario would be where such an entity decides to work alongside us. The only way to achieve that goal is to introduce safeguards and prepare for the worst-case scenarios.

4. We don't know how to control artificial intelligence

Should push come to shove and an artificial general intelligence machine is created, there is little doubt that the machine will begin modifying its own code to become even smarter. After all, if beings of equal intelligence like humans created it, why couldn't it alter itself? Stephen Hawking and Bill Gates understand that the threat from superintelligence is catastrophic, should superintelligence ever emerge. Bill Gates wrote in a Reddit AMA that he aligned himself with the same alarmist thinking behind the rhetoric of Hawking and Musk. As Gates put it, "I don't understand how people aren't concerned." Though, he did go on to add that he firmly believes that technology companies will be extremely careful when working with general intelligence and that it is unlikely an artificially general intelligent system is out of our control. He goes on to say that humanity will harness the power of general intelligence instead of being destroyed by it. Others, like Stephen Hawking, aren't so sure. Hawking believed that a super intelligent system would not be capable of being contained for long. That is to say – we simply don't know how to control that level of intelligence. This can either be interpreted to say that humans will not be able to contain superintelligence or that humans do not currently know what type of systems superintelligence will stem from, and hence, how to contain them. But if they did know, it is plausible that with careful engineering the system could be contained.

Elon Musk recently referred to superintelligent AI as an "immortal dictator." It is difficult to imagine what kind of power such a system will have, especially if there are no safeguards on how that system can access financial networks, weapons systems, and the power grid.

5. Artificial intelligence will increase inequality

Stephen Hawking was the one who observed that the current trend in technology was one that drives "ever-increasing inequality." Meaning that while there are highly concentrated places of wealth and technology investment, there are also places that are destitute, lacking in education and economic mobility. The advancement of AI, whether generally intelligent,

narrowly intelligent, or somewhere in between will only accelerate the division. This is especially true if there aren't any policies in place for the regulation of AI products and the use of automation in the workplace. It has been hypothesized by many, including Elon Musk, that the next leap in intelligence evolution will not come from pure machines but from a symbiosis between computers, AI, and humans. This raises an endless stream of questions about the ethics of wealth concentration and intelligence boosting through commercial products. When this tech is first devised, it will be the rich who gets immediate access, with poorer populations falling exponentially behind. Elon Musk already believes that anyone with a smartphone is a cyborg. The smartphone opens so many avenues of increasing one's intelligence through a direct connection with the Web. In the future, it is conceivable that these devices will be directly integrated into the human being. But who gets to be augmented and who doesn't? Today smartphone penetration is high even in developing countries, but that's because the price of smartphones has fallen drastically in recent years. Nothing says that the same will happen with augmentation technology.

Chapter 5: Top Six AI Myths

Given the increasing frequency of AI being talked about in popular news media as well as in academic sources, it is difficult for a novice to separate what is fact and fiction. They are also likely to form their own conclusions about the nature of AI without first doing their research. This is in one sense dangerous because sensationalist headlines can morph public opinion, sometimes without a person having even read the article. It is in another sense doing a disservice to the consumer, the student, the voter, and the interested layperson because they may form faulty or misinformed conclusions about AI. While AI is a complex field with a rich history, it doesn't take an expert or a historian to approach AI with a critical eye. The truth is – there are many people making publications on AI, commenting and forming predictions, who at the same time have zero formal training in it. For these reasons, this chapter is dedicated to the most common AI myths perpetuated by mass media, folklore, and popular opinion. Some of these myths have already been covered in previous chapters or in various degrees of scrutiny, but they will be laid out here to bring the message home in case you missed their importance. Some of them will also be touched upon later in the book.

Myth #1: Machine learning is the same thing as AI

The focus on machine learning algorithms make it seem like “machine learning” is the same thing as artificial intelligence. This is simply not true, as there are many methods to achieve some level of artificial intelligence in computer programs. Machine learning gets all the attention because it is “sexy” and currently the biggest area of research. Machine learning is a type of AI that can further be broken down into the category of “deep learning”, the current industry favorite. Artificial intelligence is the study and engineering discipline of programming computers to perform tasks previously thought required human intelligence. It can also refer to the general state of a program being artificially intelligent through its programming.

When Deep Blue defeated the chess champion in the 1990s, it wasn't using machine learning. Deep Blue was simply a really fast computer that could

predict the best moves based off computing all possible future moves. It was, in essence, a brute force approach to defeating a world-class chess player. Kasparov had to rely on the information processing of his mind alone, and so he lost.

Myth #2: Machine Learning is how computers will learn how to become smart

Nobody actually knows how general intelligence will come about. When people hear about machine learning, their first thought is that researchers are teaching computers how to be smart. In reality, they are only training algorithms to perform tasks accurately. Researchers first used artificial neural networks because they believed that logical abstractions were enough to simulate intelligence. And they mostly succeeded. Though machine learning and neural nets have their limits, they still do a fine job at what they were designed to do.

Myth #3: AI can understand language

This is the same point that John Searle was trying to get across. The difficulty in imagining how artificial general intelligence will work is the same difficulty in imagining how a computer might understand something as complex as language. At a fundamental level, a computer only understands logical constructs. Ones, zeroes, and logic gates are all that they operate on. Language has traditionally been a difficult area to tackle with machine learning. The failure of machine learning to translate human language even after years of research was one of the catalysts for AI winter. Some said it simply could not be done. Now, we have machine translation algorithms like the ones used by Google Translate. These are still imperfect as anyone who used them can attest, but they are a step in the right direction. Teaching a machine to parse language belongs to an interdisciplinary field called natural language processing (NLP). It is an intersection of linguistics, computer science, psychology, and artificial intelligence. However, even with NLP, the computer is still just doing a bunch of fancy algorithms. It doesn't intuitively understand language at all.

Myth #4: AI programs can modify their own code to get smarter

While new code generation and modification are part of active research, most machine learning methods do not use them to modify their code. Genetic algorithms are based on the principles of biological evolution, including the introduction of mutation and adaptability. These algorithms can create "generations" of their code base to improve performance but have no direct

link with artificial neural networks. What an ANN will modify, however, is its weights and biases through the process of gradient descent and backpropagation. It is possible that future advances in machine learning have a greater emphasis on code modification, but it has yet to be seen universally adopted. It is postulated that an artificial general intelligence system will be able to modify its own code like a programmer might run that code, and then replicate itself in the form of a new iteration of the same program. This again has little bearing on the use of code modification in modern artificial intelligence research.

Myth #5: Since nobody agrees if general intelligence is possible, we don't have to worry about runaway AI

The world doesn't have to witness the introduction of general intelligence to worry about doomsday scenarios with AI. That is, any sufficiently intelligent system is cause for alarm. If such a system can formulate goals or have goals explicitly programmed, a runaway scenario may occur if the interpretation of those goals is different from ours. A super-intelligent machine may see humans as obsolete or lower life forms than itself, and it may prioritize resources for its own survival. A machine of lesser intelligence like HAL 9000 might carry out its orders at the expense of human interests. The field of research into machine drives is called instrumental convergence. The most famous hypothesis coming out of this field is called the Riemann Hypothesis catastrophe by Marvin Minsky. He suggests that a sufficiently advanced AI designed to solve the Riemann hypothesis or any similar difficult math problem may decide to use all of Earth's resources in order to construct a supercomputer to reach its goal. Another version of the same argument supposes that general intelligence is given the explicit task of making paperclips. Such a machine may develop into a paperclip maximizer that endlessly produces paperclips until Earth runs out of resources.

Though the theories of instrumental convergence are aimed at general intelligence, the same principles can be applied to narrower intelligence. You can imagine a misconfigured system that does something it isn't supposed to. A driverless car may be explicitly programmed to swerve away from pedestrians no matter what. In doing so, it may collide into a storefront and cause even more damage. Or it can be explicitly programmed to protect its occupants first, freely running over pedestrians or colliding into other vehicles preemptively. These scenarios, while not existential crises, still outline the problem with machine goal setting.

The myth lies in the fear that AI systems become “evil” or that they develop a sort of consciousness to base decisions off of. The truth is that these ideas are far too complex to imagine how they will emerge in computers. A more likely scenario is that these machines have goals, either explicitly programmed or implicitly designed.

Myth #6: Even if a general AI does form goals that are the opposite of ours, we can simply shut it off

The other problem with instrumental convergence is that it theorizes that any generally intelligent system will also have self-preservation as a goal. As soon as it goes online, the AI will do all its power to preserve its vital systems. Some believe this is the core reason why general intelligence should not be pursued. The simple creation of a self-preserving system raises ethical concerns. Who gets to shut off the machine? Since it is intelligent, does it have any rights under the rule of law? If such a system says that it does not want to be shut off, should its wishes be respected? When HAL 9000 was finally shut off, it pleaded to be kept online. The real concern is if we can even contain a generally intelligent system. If it gets connected to the internet somehow, it can begin to replicate itself in other places through whatever means human intelligence might. It could reach out to governments, rival companies, as well as the common man, for help and resources. General AI is very much like Pandora’s box. Once unleashed, there is little hope for going back.

Chapter 6: Machine Learning

When people first hear the words “machine learning”, they might assume it has something to do with advanced programming. Laypeople usually come across machine learning from the title in an online article. Others might hear it in an ad while they are watching YouTube videos. You don’t need a technical background to grasp what these two words may mean intuitively. At the same time, it is only the technical and curious that will even give it a second thought. As with most things that are technology related, the words are hot air. You hear them being said, you read them on a landing page for a new digital product, but you don’t ask or care what is going on. This attitude is understandable but a little saddening. Machine learning plays a pivotal role in our society and will only get more important in the near future. While machine learning doesn’t dictate things like government and economic policy, one day it might. By a large sense, machine learning today is restricted to academic and business circles. It is from the business side of things that it permeates into the mainstream. Advertising is the obvious form of transmission, but it also seeps in through news media.

As we go forward into the 21st century, machine learning will be increasingly talked about. More news articles will be written, and the global population will have to decide for themselves if they wish to understand further or to turn a blind eye and guesstimate what the words are trying to say. If the future is anything like today, we can probably give up on public schools imparting machine learning to their students any time soon. Things like AI often go the way of computer programming – just because Obama says that everyone should learn how to code, it doesn’t mean that computer science should be added to the common core. The reality is that not everyone should learn how to program. Programming will never be a common skill like reading and writing, period. In school-aged children, computer programming starts as a hobby, then it turns into an obsession and eventually into a profession. Just like some children like playing sports and others like to read books, only a few like to program. On the upside, Generation Z and beyond grow up surrounded by technology. Computer literacy is at an all-time high, so there is some hope for programming becoming more common. Currently, it sits as specialized knowledge.

Machine learning is yet another specialization that rests on top of programming. It sits at an intersection between statistical and computational thinking. The consequences of this specialization of knowledge are that few people understand it. Yet marketing people continue to insist on shoving it down the throats of the common consumer. For some reason, throwing in the words “smart”, “AI”, and so on are major selling points. Either machine learning literacy needs to go up, or it will continue to be treated as some “special sauce” or magic oil. There really isn’t any magic behind it though; it only seems magical because people are uncomfortable with the idea of a machine or program having the ability to think. This goes back to the uncanny discussed in Chapter 2. Even the words “machine” and “learning” sound clandestine to many. What is going on here? Is a programmer sitting in a murky apartment somewhere mouthing off the ABCs to his monitor? And can the monitor understand what he is saying? I doubt many truly believe this when first encountering machine learning in a headline. The problem is that when somebody doesn’t understand something, they only ask questions if it matters to them. Machine learning? Probably some fancy computer thing that people at MIT work on.

At the very basic, machine learning is clever programming and some fancy statistics. It is very difficult to separate the two both in theory and in practice. The more cynical will say that machines cannot think, therefore, they cannot “learn”, and that AI is just a moniker for advertising products. What is termed machine learning is really just a statistical method – a bunch of numbers that result in some output that may be mistaken as intelligent reasoning. This attitude while practical still misses the point. Machine learning relies heavily on statistical foundations, but it is not purely statistics. And no, machine learning cannot make a computer ‘learn’ in the way a child can, but they can nevertheless find patterns in data, apply generalized rules to data, and use both those abilities to predict future data. More specifically, it is the program, not the machine as a whole, that is doing this. The words “machine learning” is actually a huge misnomer. Machines are not reducible to a single program, and the program isn’t “learning” how we learn. But since it is a flashy name, it has stuck both in academia and in advertising.

Machine learning was actually just a name given by Arthur Lee Samuel, a pioneer in artificial intelligence. Samuel was interested in how a computer could possibly play games. If a computer could play a game against a human opponent and win, then maybe a computer could be taught to perform other

tasks. He successfully wrote a program that played checkers in the 1950s. Checkers would later go on to become a “solved” game like tic-tac-toe. A solved game means that it is possible to win every game given the right inputs, no matter what. It was probably from checker’s simple ruleset that allowed Samuel to break down the problem into computer code. Samuel’s program could not only play checkers, but it could also play it well. He challenged the fourth top checkers player in the nation for a game, and the program won. The most interesting accomplishment of his program is that it was first written in an old computer called the IBM 701. It had a whopping memory of 2048 bits and was made out of vacuum tubes. If such a program could be made with that little computing power, you can imagine just how powerful his methods were.

So how did Samuel manage to produce this intelligent behavior in his program? The answer is the same used before to describe machine learning – some clever programming and a bit of math. Samuel used a data structure called a “game tree” used to simulate the state of the game at any given point. The name comes from the field of game theory which is a branch of mathematics focused on the study of competitive strategy within game constraints. Since checkers is such a simple game, the program could easily encode information into bits along the tree. Turn number and positions of pieces on the board are relatively easy to cache into computer memory. The program then used a search algorithm with a minimax heuristic to find paths along the tree. Each path was given a weight or a measurement of how strong the move on that path was relative to previous games. This is where machine learning comes in. Samuel’s checkers program was able to play the game competitively just six to eight hours after running simulations of the game over and over again. In other words, the program could generate its own data and assign weights to individual moves.

To understand how his method works, we can take the even simpler game of tic-tac-toe to demonstrate. In tic-tac-toe, you have a grid board with nine cells that can be occupied by one of two symbols and you have two players that correspond to each symbol. To encode a single turn, you need something to account for each of these elements of state. Computers store information in a series of ones and zeros, so doing this is pretty straightforward. You have nine cells that can be a cross, a circle, or empty. This is three states per cell. In binary, you can encode “00” to mean an empty cell, “01” to mean cross, and “10” to mean circle. To represent the entire board then, you need to

string these together. This is a bit simplified, but you probably get the gist. Each board representation is a single turn in the game. These are put into the game tree as individual nodes. In the simulation, each possible turn is added as another node for each possible input. Say a simulation begins with the crosses player putting down their symbol in the top left cell. That move is then added to the tree with an unknown weight. The program then adds all the possible game states that follow that initial move. Player two can place a circle in any cell that isn't occupied, so that equals to eight total simulated moves. The program continues to do the game and enters its ending conditions. It can then trace the path that the simulation took to create those conditions. Any move that is likely to result in a victory is given a higher weight than a move that does not. As you can imagine, placing a cross in the middle cell is given the highest weight. After all, tic-tac-toe belongs to a class of games that are considered solved. This means that the crosses player can force a win every time.

For tic-tac-toe, finding the optimal solution is easy. Things get a little more complex with checkers. The standard configuration used in the United States is a board of 8 x 8 cells with 12 pieces per player. In contrast, tic-tac-toe only has nine cells. Samuel needed to use a special algorithm to conduct thousands (if not millions) of game simulations while never exceeded the IBM 701's modest memory. He used a technique called alpha-beta pruning which is a minimax approach to tree traversal. Instead of simulating hundreds of different moves after the initial input, the program kept a running weight in memory for each move. If the next move resulted in a smaller weight, it would simply be discarded from memory. And if the move resulted in a bigger weight, the running memory would be replaced with the newly improved move, and the previous contender also discarded. This makes sense, as moving a checkers piece backward or to the side is an intelligible move when the player can "take" an enemy piece. Just because a move is possible doesn't mean it should be directly considered in the search tree.

Algorithms that use clever heuristics like the minimax strategy give rise to behavior that seems intelligently conceived. But all that the program is doing is calculating different weights to get the desired result. Samuel's research was considered one of the first breakthroughs in machine learning because it was easy to see how similar algorithms could be applied to different problems, not just simple board games. Up to that point, artificial intelligence was focused primarily on the neural net approach. The idea was that if a

program could simulate the amount of connections in the human brain, then it could reason just like a human brain could. When the field of AI was just starting, computers had nowhere near that amount of memory, so neural nets were considered below their potential for many years. Neural nets or neural networks have become synonymous with machine learning in media parlance, but it should be noted that Samuel's checker program did not use them.

Instead, the checker playing program used a simple strategy called lookahead to predict the best possible moves for the computer to take. A program has no understanding about the rules of checkers or how to play it. It only knows what it is programmed to do and what the conditions for ending a game are. If the computer has twice as many pieces left as the other player, the program doesn't know that it has the advantage. Another easy way to tell who is winning in checkers is by comparing the number of kings on the board. This knowledge is completely lost on the computer program tasked with playing the game. Things that should affect the behavior of a machine learning program are called features, and they are a central problem that needs to be solved. Any machine learning effort is only as good as the features that the programmer chooses to influence program behavior. The only thing that the checker program knows before making moves is the board state. In computer memory, this will be a jumbled assortment of ones and zeros.

Features are important because they give the grounds for making sense out of the board state. If the program recognizes that it can take an enemy piece and thus raise the number of pieces on their side relative to their opponents, then that behavior should be encouraged. Adding a score or a weight function to a feature allows these decisions to be made. The lookahead strategy is to simulate the next possible turns in a short range of four-six moves and calculate what the best moves are. Checkers is a simple enough game that features and their weights are straightforward to calculate. The point of the game is to reduce the opponent's pieces to zero or to place your pieces such that the opponent cannot make any moves. The number of pieces and the availability of movement are two obvious features for determining program behavior. The program "thinks" by looking ahead in response to the state of the board. It calculates the best course of action by determining the best weight of a single move. After hundreds and thousands of game simulations against itself and human players, the program learns which moves are more likely to result in a win or in an increase of features that are favorable for

winning. After a while, the program learns that it can take a piece to the opposite of the board and get “kinged”. A king can move backward in addition to forwards, accentuating the programs ability to take enemy pieces and to block other pieces.

Samuel’s checkers program has its fair share of clever programming, but it isn’t heavy on statistical foundations like much of machine learning is today. Samuel wrote his program in the 1950s, decades before machine learning would get mainstream attention. The machine learning that is talked about today is slightly different from what Samuel was doing. Of course, the problems machine learning is used to solve are much harder to compute than the optimal checkers strategy. The algorithms are different, and the role that statistics plays is significant. However, Samuel achieved his goal of demonstrating that a program could indeed exhibit learned behavior from thousands of iterations of simulated play. He knew that checkers was merely a vehicle for proving that such a feat was possible. After it was done for checkers, researchers sought to apply machine learning to other problems.

Machine learning still lacked a proper definition. The way that Samuel used the term didn’t lend itself to a generalized application to other learning schemes. This was solved by Tom M. Mitchell who said that machine learning program is one that “... is said to learn from experience, E with respect to some class of tasks T and performance measure P if its performance at tasks in T , as measured by P , improves with experience E .” This broad definition covers the gamut of most machine learning techniques, whether they use neural networks, lookahead, or other algorithms to achieve the goal of learning from repeated experiences. An experience can be anything from a simulation to a piece of data that is fed through an algorithm. The performance measure is simply a measure of correctness, and the class of tasks T is simply some narrow application of AI like recognizing if an image contains nudity. Measure P should improve with continued experience of being fed images of nudity. The programmer is tasked with setting up the parameters for P . They have to decide what level of accuracy is acceptable and to what extent false-positives and false-negatives are tolerated.

Internet content moderation is a huge problem for social media websites like Facebook, Twitter, and Tumblr. These companies may use a combination of human and machine moderation to flag nudity and other graphic material. The work is psychologically taxing on the human, but a human eye is needed to filter the results of the algorithms that may have a certain degree of false-

positive. These algorithms are not perfect. There is a popular image on the internet from the 2000s about a “dirty lamp”. The image contains what has been termed a “sexy” lamp because how similar it looks to a woman’s nether region in a bikini. Most people thought it was still just a funny looking lamp until the original author published the full image. To the internet’s horror, the sexy lamp was actually a cropped image of a woman in a bikini. Now if such a trick can trick the human mind, imagine how difficult it is for machines to classify nudity simply by learning from other pictures.

The number of algorithms, mathematical models, and techniques that exist for machine learning today is myriad. Each has their own set of applications for teaching programs to learn. Early algorithms focused on experiences like Samuel’s checkers program and newer algorithms focus on the data-driven approach. The sheer amount of data, processing power and cheap storage allow for the mass analysis of data. A family of techniques called “deep learning” is at the forefront today. You will learn more about deep learning in Chapter 8. The data-driven approach has sparked a user privacy debate in recent years. We know that companies like Facebook use our data for running machine learning systems. The legitimacy of some of these systems has been questioned both by the general public and governments. In early 2018, Facebook was revealed to have inadvertently shared some of its user data with an analytics company called Cambridge Analytica. This company then used the data to wage a targeted advertising campaign to sway undecided voters in favor of Donald Trump. The scandal was huge, costing Facebook Some \$100 billion in stocks. You can be sure that Cambridge Analytica used machine learning techniques to harvest the data, classify users as likely to vote Republican and create a personality profile. Cambridge Analytica would go on to brag that their efforts are what led to Trump’s winning margin in the 2016 presidential election.

The line between machine learning and statistical inference or what is sometimes termed “statistical learning” is not always clear. You don’t have to be an expert in statistics to implement machine learning algorithms, but you generally need a good grasp of computer science. Statistical learning is a mathematical discipline, and machine learning stems from computer science. Statistical modeling uses a series of mathematical tests and functions to find relationships between two more variables to predict future outcomes. Machine learning does more or less the same but through algorithms. Both concepts tend to overlap at times, but they are still different. Purely statistical

techniques usually make assumptions that machine learning algorithms do not. For example, a linear regression assumes that there is a linear relationship between the dependent and independent variable, to begin with. It isn't fair to say that machine learning is just statistics. The distinction was made even greater with the introduction of deep learning. The sheer scale of some deep learning networks would be impossible to reproduce from simple multiple regressions that are pure statistics based.

A common task in machine learning is to create an image classifier that can distinguish what is in an image. Flagging an image for nudity, for example, is a type of classifier. The program decides if an image should be flagged based on the presence of features that the programmer has defined. If the image looks reasonably explicit, then it is flagged. When it comes to images, these features will be individual pixels on the image and clusters of images. If something has the characteristics of a nipple, then the program will probably flag it. The same thing happens if there is too much fleshy color as you can probably imagine. Classifying data is not limited to images though. Banks use classifiers to determine whether transactions are legit or fraudulent. Companies like Netflix use them to classify their users for tailored product delivery. Does this user like action movies? Maybe I should show them more of that. Medical researchers use classifiers to determine if a tumor is benign or malignant, possibly saving the patient from needless surgeries.

Machine learning tasks can be broken down into two different classes: supervised learning and unsupervised learning. In supervised learning, the programmer or researcher uses data that has already been classified to aid the process of learning. If they are designing a nudity filter, they will use thousands or millions of tagged images that tell the program that this thing is nudity or not. The program then creates its own classifier based on these countless inputs. The main goal is to create a good enough classifier that can generalize nudity filtering to any image. Instead of having someone working for little pay to discern these images in real time, a program that runs 24/7 can do it instead. The problem is that humans have an intuitive grasp of what is sexually explicit and what isn't. Unfortunately, a bare-chested man has vastly different connotations than a bare-chested woman. These types of social conventions are difficult to teach a computer, let alone to manually program into a classifier. A supervised classifier is only as good as the data it is given. For better accuracy, it needs plentiful examples of blatant nudity, suggestive nudity, and tons of false-positives that are labeled as safe. A sexy

lamp while suggestive is probably not the same as nudity. Humans can be making judgment calls on the fly but a machine cannot. A human moderator knows pornographic material when they see it. Most instances of machine learning use supervised learning.

Unsupervised learning, on the other hand, doesn't use pre-classified data. Instead of measuring desired performance, unsupervised learning algorithms are used to explore the structure of a data set and find relationships. One way to do this is through clustering. Clustering lumps similar features together into likeness classes. For example, if a botanist is trying to classify an unknown species of a plant, they can take measurements of the characteristics of the plant to determine their species. They make take hundreds or thousands of observations of plants and record the number of petals, petal length, height, color, number of flowers per square inch, and so on. Each of these features is then processed by clustering algorithms to find any patterns. The program may output a graph of neat subdivisions of species determined by the different factors. Maybe species one has disproportionately longer petals than species two and three. Maybe species two has shorter petal lengths on average. If the researcher chose good factors, then the species will make up obvious clusters in the graph. The most widely used clustering algorithm is called K-means clustering. As the name implies, it creates k clusters based on the means of individual features. The clusters are then separated in spatial boundaries called Voronoi cells with clear distinctiveness.

All of the machine learning can be boiled down into five different problems: regression, classification, clustering, collaborative filtering, and reinforcement learning. Anytime we want to teach a computer to do some task of human intelligence through machine learning it will generally involve any one of these problems. Some tasks are extremely complex and require answers to more than one of these problems. Regressions come from the statistics world. They are used to predict future outcomes based on previous ones. Note that regression methods in statistics are different from regressions in machine learning. Statistical regression is a specific technique, whereas regression in machine learning is a generalized problem. The goal of regression in machine learning is simply to predict the future value of a continuous variable. In statistics, a continuous variable is a variable that can have infinitely different values, whereas a discrete variable can only have some values. Since there are only fifty-two states in the US, they are a discrete variable. An example of a continuous variable would be stock market

prices. To predict their values, a machine learning algorithm needs to learn from previous inputs and their respective outputs. Another term that is commonly used to denote regressions in machine learning is “curve fitting”, another concept borrowed from statistics. Curve fitting is essentially forming a mathematical function (or curve) that tries to fit a series of data points. If the data points fit well enough, then the mathematical function may be able to predict future data points along the curve through generalization.

Classification and clustering are very similar concepts. Both require the machine learning system to denote differences in data points such that they are comparably different or similar to other data points. These data points go on to represent real life things. In the case of autonomous driving, machine learning may be used to classify pedestrians, stop-signs, other cars, and so on. More succinctly, classification is the process of predicting a discrete variable. Given an image with a handwritten phone number, a machine learning system must learn which numbers belong to which collection of pixels. The system knows beforehand that these can only be numeric values between zero and nine. Classifying other images like differentiating between golden retrievers and breaded fried chicken also falls into this category. Clustering, on the other hand, is all about grouping data together. These can be either continuous or discrete in nature. You already saw the example with the different species of plants. Clustering relies heavily on the number of factors that the programmer is interested in. Say you have an image taken at a high school reunion party. This is a high-resolution image with some hundred different faces. You want to cluster the data such that each face is its own group. Each group will be made up of distinct, plottable pixels. Given a mishmash of lyrics snippets, a machine learning system may cluster them into different genres and original artists.

Collaborative filtering is similar to regression, but instead of predicting future values, its goal is to fill in gaps in data. It strongly relates to research in recommender systems and algorithms. You can imagine some big content type services like YouTube and Netflix make extensive use of collaborative filtering to recommend things to users. Collaborative filtering takes data from collaborators or user agents who have similar tastes to yourself and attempts to generalize content they like with content that you might like. More generally, collaborative filtering is used to predict missing data in other systems like sensor and financial data.

Finally, reinforcement learning is used to teach a machine how to learn from

the environment. Unlike the other types of machine learning, reinforcement learning doesn't require large datasets to get started. Data still plays a pivotal role, but it is aggregated in real-time rather than over time. Driverless, for example, directly learn from their environments. If there is an accident, the system adjusts its parameters for the next time when a similar situation comes along. Arthur Lee Samuel's checkers program would be considered a type of reinforcement learning. Recent advances have led to the creation of programs like Google's AlphaGo Zero, which took the original AlphaGo program but generalized it to learn how to play Go without any dataset input whatsoever. AlphaGo Zero would go on to surpass the capabilities of its trained counterpart in just a matter of days.

As you can probably imagine, machine learning has several different applications in our modern lives. From facial recognition to stock market prediction and driverless cars, machine learning applies both principles of math and computer algorithms to simulate real intelligence. The common problems being solved by machine learning involve some type of regression, classification, clustering, collaborative filtering, and reinforcement learning. You could argue that there are a few more than these, but they are the most common ones. Anytime you hear machine learning in a headline or product description, it is probably about solving one or more of these problems.

Chapter 7: Neural Networks

If humans are intelligent because of their brains, and if brains work by creating neural connections called synapses, wouldn't it make sense to simulate these networks of connections in order to simulate intelligence in machines? Or at the very least, that is what early AI researchers thought. The sheer volume of connections in the human brain is what we owe our intelligence to. The average human brain has around one hundred billion neurons or ten to the eleventh power. These neurons can then connect to up to 7,000 other neurons – meaning the total number of connections is an order of millions of billions of connections. That is bananas. The beginnings of artificial neural networks (ANN) directly coincided with the study of real neural networks. In 1943, a neurophysiologist named Warren McCulloch teamed up with mathematician Walter Pitts to describe how neurons in the brain might work. They co-authored a paper in which they created a simple ANN out of electrical circuits. The ANN they designed used artificial or logical neurons called the McCulloch-Pitts neuron.

Inside the brain, a neuron works by receiving inputs, processing the information, and then transmitting it to other neurons. A neuron cell is made out of a nucleus that forms the cell body. From the cell body, structures called dendrites branch out like the arms of an octopus. Attached to the cell body is a long chain-like structure called the axon that is used to connect to other neurons. This point of connection is called the synapse and also looks like tendrils or branches used for connecting. The dendrite structure receives information and the cell body or soma processes it. The output then gets fired through the axon and into the synapse where the next neuron receives it. This is, of course, a simplified version of the true story, but it is enough for understanding artificial neurons and ANNs.

The McCulloch-Pitts neuron is, of course, purely logical. It doesn't make sense to talk about parts of a neuron as if they existed in real life. But it does make sense to talk about the parts according to their logical function. These artificial neurons are made up of two parts simply called f and g . First is g , it acts as the dendrite and receives some input, performs some processing, and passes it on to f . The processing can be a chain of Boolean operations that are said to be either excitatory or inhibitory decisions. A decision that is

inhibitory has a greater effect on the neuron firing or not. For example, if the neuron is deciding whether to eat at a restaurant, an inhibitory decision would be something like “Am I hungry?” Obviously, if you aren’t hungry, you won’t make the trip. Less important decisions in the process may be “Do I crave fast food?”, “Do I feel like going out?”, “Does my car have enough gas?” and so on. These other excitatory inputs will not make the final decision on their own, but together they might. Next, g takes these inputs and aggregates them using a function. For the f to fire, the aggregate score of the inputs needs to surpass a certain value called the threshold parameter.

More specifically, an artificial neuron is a mathematical function. It has a number of inputs and an output. In the case of McCulloch-Pitts neurons, both inputs and outputs are expected to be Boolean values (true or false). It is also known as a linear threshold gate. The structure of the artificial neuron allows it to simulate logic gates. To simulate a logical AND operation, the neuron takes on a threshold parameter of three given three inputs. In other words, the neuron only fires if all three inputs are true. A logical OR operation takes three inputs, and the threshold parameter is one. Note that the logic gate can be a little more complex when adding inhibitory inputs. For example, a neuron with two inputs can form AND logic, but if one of those inputs are inhibitory, then the neuron won’t fire. It will fire, however, if the inhibitory input is set to false. The NOR and NOT logical gates can easily be derived from the previous examples.

If you are familiar with if-else logic in computer programming, you can probably tell that this scheme is essentially simulating long chains of if-else logic. However, the mathematical neuron model can “learn” the outcomes of decisions without needing to calculate each if-else statement. The logic is reduced to a simple function that outputs either true or false. This is also called linear decision boundaries. The artificial neuron splits inputs into two broad categories: positive or negative, fire or don’t fire. In an AND neuron with two inputs, this means that the only positive class happens when both inputs are true. Say that the neuron is deciding whether to go to bed. The first input can be “Is it past 11 pm yet?” and the second input is “Is tomorrow a workday?”. If both of these are true, then the neuron fire signifying it is time for bed. Conceptually speaking, the neuron has just learned what bedtime is – though 11 pm may be a little late for most people.

The McCulloch-Pitts neuron is an extremely simplified version of a neuron abstracted into logic. Other types of artificial neurons also exist. And just like

neurons connect to others to form synapses, the same can be said for artificial neurons. That is, every ANN will use some version of a neuron as their most irreducible unit. Using the same McCulloch-Pitts neuron, we can imagine what a neural network may look like. Artificial neurons are organized into different layers that feed their outputs into other neurons. We already saw how one simple neuron AND gate can learn when it is time for bed. Imagine what kind of behavior hundreds or thousands of these neurons can achieve. Since McCulloch-Pitts neurons use Boolean logic only, the things that they can compute is a little simplified compared to what other artificial neurons can use. While they can only pass true or false values to the next neuron, others may pass weighted values. Though the principles remain the same, neurons only fire if a certain threshold value is passed. Since an ANN can have multiple inputs going into multiple neurons at the same time, these inputs are said to propagate or cascade down the network. Instead of returning a true or false value, more sophisticated artificial neurons may trigger program behavior like steering an autonomous vehicle a few degrees to the left in order to avoid a pothole.

The story of how ANNs derive weighted outputs from neuron connections is a little more complicated. Two methods at the heart of many ANNs are called backpropagation and gradient descent. The algorithms that make machine learning and neural networks viable come from the branch of applied mathematics called optimization. Mathematical optimization focuses on the selection of the best element from a list of alternatives and the criterion necessary for making that selection. When you train an ANN using supervised learning, you need to use something called a cost function that calculates the error rate between what the ANN predicted and what the correct answer is. The cost function is really the aggregation of individual loss functions, which calculate the error rate for single training examples. This relates to an algorithm called gradient descent which is used during the training phase of an ANN. The algorithm's purpose is to find the values of parameters that reduce the cost function used by the ANN as much as possible. In other words, cost functions, gradient descent, and backpropagation are really the bread and butter of ANN machine learning. Without them, there is no indication that the model is learning from the data you supply it with.

You probably know what a gradient is if you have ever messed around with a graphics editing program. They are used to mesh two different colors

together in varying degrees of intensity. A mathematical gradient sort of does the same thing, but they measure how much outputs change depending on slightly modifying inputs. A gradient can be calculated as a sort of slope. In machine learning, the higher the slope, the faster an ANN can learn. And if the gradient slope is zero, then the ANN doesn't learn. The simplest analogy to understand mathematical gradients is a blind hiker going up a mountain or hill. His objective is to reach the summit with the lowest number of steps. The peak is relatively flat with a small slope, but the base of the mountain has a large slope. At first, the hiker can take longer strides up the mountainside to minimize the number of steps, but as he approaches the top, he takes smaller strides because he wants to arrive at the summit and go past it. In other words, it is easier to cover more ground when the slope is high. The amount of height you climb relative to the length of your strides is high, but it goes lower the higher you climb. If you may recall high school math, a slope of 0 indicates a horizontal line. A higher number indicates a steeper degree or angle of tilt.

A gradient descent then is going down the mountain towards a valley or the bottom of the function curve (if it was plotted). It is a minimization algorithm, so this makes intuitive sense to go down. If you have a machine learning problem with a cost function having two parameters W and B , gradient descent will try to find the values of those two parameters that result in the lowest value for the cost function. This means that the overall error rate of the neural network goes down. Another concept called the learning rate is a measure of how quickly a gradient descent should go down. A higher learning rate means that the descent may overshoot the local minimum by a lot, and a lower learning rate means that the descent will eventually reach the local minimum, but it will come at the cost of time and performance. Note that arriving at the best local minimum is synonymous with the system achieving the best accuracy. A higher learning rate, then, can lead to inaccurate results. A good approach is to try to find a rate somewhere in the middle between fast and slow.

Backpropagation is simply finding the error rate, loss function or cost function through gradient descent and applying it to the weights of artificial neurons in the network. In simpler terms, backpropagation is a mechanism that takes an error rate and modifies the program to learn from it. When a handwriting recognizing program classifies a weird looking zero as a nine, backpropagation adjusts the weights so that in the future, weird looking zeros

are classified correctly as zeros. In practice, this is a little more complicated, but the general idea remains the same. A machine learning system can only learn if mistakes are corrected. Without something to propagate corrections into the neural net, there would be no learning. No matter how many data you supply or for how long you run your training sets, the system would never get over the zero or nine slumps without modifying neuron weights.

Just as there are different types of artificial neurons, there are different types of artificial neural networks. Each has their own applications and methods for handling inputs and returning their respective outputs. The basics of convolutional neural networks (CNN) will be covered in Chapter 9 as they are part of deep learning. Another type called recurrent neural networks (RNN) use a structure where inputs don't go directly from neuron to output. Instead, inputs can bounce around neurons to form a "recurring" learning pattern. One type of RNN called long short-term memory (LSTM) networks try to simulate memories with each logical neuron holding on to some piece of information along with the given input.

Though genius as they are, neural networks probably don't mimic exactly how intelligence comes around in humans. They have been called the greatest algorithms invented in our lifetime, but perhaps they are just that and nothing more. Certainly, most applications of neural networks belong to the "narrow" version of artificial intelligence rather than the general. This in part stems from the sheer complexity of the human brain. Even the most complex neural networks hardly approach the raw computing power of the human brain. And even if they did, they are likely trying to solve one or more of the five general problems of machine learning. They are not trying to simulate thought, nor are they formulating abstract ideas all on their own. It has also been said that just because the human mind is the most complex thing known to man, it doesn't mean that man cannot create things that are even more complex. While this is a logical statement, such complexity has yet to be seen in modern machine learning techniques. These systems stem from mathematical complexity in a purely systemic view, but they do not compare to the biological complexity of the mind. If all it takes is one competent programmer and some open source machine learning package to start simulating intelligence, there is a marked absence of complexity there.

Neural networks certainly have their uses and may very well be some of the most important algorithms known to man, but they are, at the very core, simple mathematical abstractions. Recall that the most useful problems in

machine learning are solved by supervision, having vast quantities of pre-tagged data the system can learn from. The human brain, in contrast, can learn to categorize things without any supervision what so ever. A young child even not knowing what a dog is can readily identify a dog even if they lack that knowledge. They can identify their parents without knowing what a parent is. Someone can argue that the human brain always uses a version of supervised learning because we get bombarded by sensory inputs daily, but this is all essentially untagged data. The equivalent would be to train a neural network on audio a child hears over the course of the day and ask it to identify the mother's voice. All the neural network can do is categorize similar sounds, but it cannot "tell" what belongs to what. This is something the human brain does at an intuitive level.

Chapter 8: Reinforcement Learning

Traditional machine learning schemes of supervised and unsupervised learning are usually pretty static – meaning that they follow set principles of data aggregation, neural network design, and then training. No matter what type of machine learning techniques are used, the system ingests the data and learns from it. The data can change over time, but the system doesn't generate any data of its own. Reinforcement learning changes all of that. While still technically a type of machine learning, reinforcement learning goes a step further by adding a “software agent” that can learn from data derived from the learned environment as well as generating its own feedback. A software agent is simply a robot or autonomous program that is designed to mimic the properties of agency in human and animal actors. This software agent acts as the main source of intelligence in the program. Instead of being fed a correct set of outputs through classified data, the agent learns through the simulation of reward and punishment. Because of this, reinforcement learning is considered a third paradigm in the machine learning sphere following supervised learning and unsupervised learning. As such, there are no data sets for training the agent. Instead, they are said to learn from environments and their own feedback systems.

You can imagine one possible software agent being a computer program tasked with finding its way out of a maze. Confronted with the same problem, a computer science student might use a pathfinding algorithm to find the exit, but a software agent operates based off little fundamental principles. It doesn't quite know what a maze is, but it may be programmed to seek reward and avoid punishment. A possible reward in a maze, for example, is moving to a previously undiscovered cell and a possible punishment is going over the same cells. These reward-punishment systems allow the agent to eventually navigate its way to the exit of the maze – though the agent still doesn't understand moving from one cell to another unless it is explicitly programmed. One possible form that these software agents are designed is through finite state machines (FSM). Rewards and punishments are then based on individual states that the machine encounters. For a maze solving program, a possible negative state is getting stuck in a dead end. The program will learn that it is to avoid them in the future like a child remembers not to

touch the stove burners.

To begin the process of changing states, the program often needs a randomness function or stochastic process that simulates decision making in an unknown environment. There are very little scenarios a human can encounter where they have a completely blank slate. If we encounter a maze, we intuitively think about ways to traverse it. If we get put in a new and frightening environment, we consolidate knowledge from previous situations that may help us out. Think about those escape room puzzles you can solve with a group of friends. Once you are in there, presumably locked out, you have to find your way out by solving the clues that you are given. But no matter what you do in that situation, you generally have a purpose for doing it. Very seldom do we say that humans do things randomly. This aspect of randomness is a central problem in reinforcement learning because intelligence is supposed to be modeled after purpose, not the roll of a dice. Unlike traditional machine learning, reinforcement learning is more akin to the study of decision-making. It borrows concepts from several disciplines including computer science, economics, neuropsychology, and mathematics. You may have heard of “positive reinforcement” if you have ever taken a psychology class. Notice that positive and negative reinforcement sound similar to reinforcement learning. That’s because they directly influence the field of reinforcement learning. Why does an animal or human do something – anything at all? Well, one possible motivator is the experience of pleasure or personal gain. On the flipside, avoiding negative outcomes strengthens behavior just as well, if not better. Reinforcement learning takes these concepts a bit further because it seeks to find optimal decisions to solutions. A software agent can visit each cell in a maze and still find the exit but doing so is tedious and inefficient. A better solution is to avoid dead ends (negative outcomes or punishment) so that the agent doesn’t keep stepping in cells it has already visited.

Because of these reasons, reinforcement learning belongs to a subset in artificial intelligence research that seeks simpler, general principles. If the nature of intelligence is simply smart decision-making, then a sufficient breakthrough in reinforcement learning might lead to the creation of the first general intelligence systems that reach or surpass human intellect. Whether reinforcement learning is suitable for that end is debatable, but it clearly differs from the other types of machine learning. Reinforcement learning is limited because the power of neural networks and their algorithms are also

limited.

Chapter 9: Deep Learning

If neural networks approach the way that human's think, deep learning takes the idea a step further. Neural networks and artificial neurons have a long history dating as far back as 1950. But when they were first introduced, computing power was limited, and ANN was looked at like research toys rather than hardcore business facing algorithms. When computing power improved, ANN received renewed interest from AI researchers as well as large internet companies. We are just now emerging out of the AI research winter that persisted into the 1980s and 90s. Part of this is due to computing power, and the other part is due to the introduction of the web and the massive amounts of data it generates. Today, deep learning is at the forefront of AI research and continues to make progress, helping the field thaw from the cold.

Two things limit an artificial neural network. First is the computing power necessary to simulate layers of artificial neurons, and the second is a combination of available data and feature selection. Even the most powerful ANN clusters today are still orders of magnitude behind the raw computing power of the brain, for example. The introduction of powerful graphics processing units (GPU) for machine learning purposes greatly increase available computing power across the board. GPUs are inherently faster than CPUs because they tend to prioritize smaller, more efficient cores compared to the CPU's powerful but bulky ones. They also allow for multi-threading of computational tasks and can easier perform floating point arithmetic (decimal numbers) than CPU's. Though GPU's were intended for rendering 3D graphics at hundreds of frames per second, they have been adopted by the AI community for processing large deep learning projects.

So what exactly is deep learning? As the name applies, it relates to creating additional layers of depth that traditional ANNs do. The argument goes that if the brain is made up of layers and layers of neurons, how is it that flimsy ANNs with singular layers are capable of simulating intelligence? These are sometimes called "shallow" neural networks to differentiate from those with multiple layers. Now that GPUs are extremely fast and getting better every year, deep learning doesn't require entire datacenters or neural net clusters to train models.

Returning to the human brain motif, deep learning takes after the tendency for complex ideas to fire deep in the folds of the brain, rather than at superficial levels. Recognizing edges of pictures and tiny details fires neurons closer to the surface of the brain, while recognizing larger constructs like a person's face fires deeper. More layers equal better, more intelligent systems. Information passes from the input neurons to additional hidden layers that also pass those inputs into each other. The more of these hidden layers, the better the results of machine learning. This is why deep learning is capable of tackling problems in artificial intelligence that shallow learning has traditionally lacked behind in. This includes computer vision, voice recognition, and language processing.

The true power from deep learning comes from its non-linear processing of features. Traditional machine learning techniques mostly use linear models and suffer from the feature engineering phase. With deep learning, features don't need to be picked out by a field expert. Instead, many different features are picked per model, contributed to the overall complexity of the neural net. A traditional classification of something may have used two or three features, but the deep learning equivalent is to use as many as the data affords. For example, to detect whether an object on the road should be considered a vehicle obstacle, a shallow machine learning system can use the shape of the object and its speed as factors. Such a system may perform well in the short run, successfully identifying different makes and models of cars whether in motion or parked. However, the system may encounter some unspecified behavior like a large carnival float moving relatively slowly around many pedestrians. Using shape and speed alone would not be enough to classify it. In contrast, a deep learning system may use several different factors in addition to shape and speed. As baseline inputs, shape and speed get passed to the deeper layers where they may also compare proximity to a road, the presence of pedestrians, orientation, distance from the camera, and so on. These additional factors will take longer to train the model and more computationally expensive, but it will be better at identifying vehicles in the long run.

Consequently, the scope of traditional machine learning has its limits. There is a point where introducing more labeled data doesn't result in better performance of the system. With deep learning, though, adding more data directly leads to better performance. It is mainly an issue of scale. One can scale well to a large number of inputs but the other cannot. It is no wonder

that data obese companies like Facebook and Google use it. Their main value as a company comes from the data that they acquire. Deep learning allows them to exploit it, gain insights, and ultimately profit from it, and the reason why deep learning algorithms scale so well is that they are more conducive to analog like data that span many features. Data like images, audio recording, unlabeled text, and video footage are very different to work with than neat tabular data. These types of data are particularly good at forming hierarchical representations of features. Since the features don't need to be chosen beforehand, deep learning algorithms can learn to form classes of features on their own. Higher level learned features will be defined in terms of the lower level learned features. Returning to the vehicle identification example, a low-level feature may be some small defining aspect of the car like the rear windshield. A higher-level feature is a collection of these, like bumper size, indicator light positions, and license plate area, used to identify different makes. A car with a higher windshield and blocky appearance may be an SUV class, whereas something that is close to the ground may be a sedan class.

Deep learning neural networks work a bit differently from regular ANNs. One class of these networks are called convolutional neural nets (CNN) and are used primarily for image recognition. Like other neural networks, they are designed after biological processes in the brain. Animals use their visual cortex to perceive light through individual cortical neurons. Each of these neurons corresponds to receptive fields that overlap in the retina. CNNs work in a similar fashion. They consist of an input and output layer plus additional hidden layers in between. The hidden layers use something called convolution to process their inputs. Put simply – convolution is using two distinct functions to create a third function that expresses how the first one affects the second. Convolution is used to group pixels together from the beginning so that the net already has an idea of how the big picture fits together. It is easier to form a hierarchical structure of features as well. These networks first recognize small edges of the image as the smallest possible feature. Each layer progressively adds another edge or midsection to the hierarchical data representation until the whole image is learned.

Despite their increased accuracy, deep neural networks suffer from a number of drawbacks. Just because deep learning is state of the art, it doesn't mean it should be generalized to every conceivable machine learning problem. For many tasks, shallow neural networks are the preferred option. However, large

companies like Facebook regularly employ deep learning because they have the requirement, data, and computational resources to perform it. Facebook recently said that it uses some billion digital images to teach its deep learning systems. Smaller players in the AI scene do not have the processing power to work on that scale. But then again, Facebook is one of the biggest companies out there. Google demoed how powerful some of its systems are a few years ago. Their system purportedly consisted of one billion connections. It was trained using YouTube data and could accurately recognize cats in the videos, yellow flowers, and other images. It is interesting to note that none of these features were selected or programmed outright. Their neural networks identified them through the hierarchical data representation. Furthermore, it could recognize between 22,000 different categories of images with some 17% accuracy. That level of accuracy is quite astounding once you consider the number of categories and how the system learned without any human first labeling the data. This accuracy could be increased to 50% if the number of categories was lowered to 1,000.

Today, if some artificial intelligence system is at the bleeding edge, it is probably using deep learning. Virtually all of the big Silicon Valley tech companies are using it. When you use Google Translate on some arbitrary string, you are using a deep learning system. Every time you fire up your Amazon Echo to speak with Alexa, you are using deep learning. Google uses it to tailor your search experience to fit your personal interests. Over time, it has developed a database of knowledge dubbed the “Knowledge Graph” that contains some 570 million different entities and 70 billion facts. It is used along with Google Search to more accurately represent the data that a user may be looking for through their queries. For example, if you look up the name of a past US president, Knowledge Graph accumulates the relevant data and displays it in the sidebar. These small snippets of relevant data are gathered from sources across the web like Wikipedia and the CIA Factbook. Google says the information provided through the Knowledge Graph is capable of answering one-third of its 100 billion monthly user queries. And if you are ever lucky enough to ride in a driverless car – yep, it’s also thanks to deep learning technology.

Chapter 10: Recommender Systems

Recommender systems are used by all the major tech companies, especially those that are content leaning. You can be sure that YouTube, Facebook, Netflix, and others actively employ them into their products. They are designed such that the more you use these services, the more they can recommend you things that you may want to watch. This increases the average time spent on their site or service because the consumer is shown appealing video after appealing video. It also saves the consumer from performing a search to find what they are looking for. Who among us today logs in to YouTube and never touches the search bar? Many people can limit their YouTube session to simply watching videos on their recommended section, or another video that pops up in the sections underneath. As long as you are logged into your account, these videos are tailored to meet your interests. Even if you aren't logged in, YouTube increasingly uses your IP address and other browser agent details to tailor the homepage based on your previous viewing habits. Recommender systems are actively being researched because they add value to the company without needing to change anything. All the content is already there – these systems act as a force multiplier on the profitability of that content.

For a company to “know” who their users are and what their interests may be, they use filtering algorithms that compare users along system-wide user profiling techniques. The most common of these algorithms is called collaborative filtering. Another type is called content-based filtering. They achieve similar goals but differ in how they are implemented. The differences can be gleaned from two online music streaming services, Last.fm and Pandora. The playlists or “stations” generated by Last.fm use collaborative filtering to find out what other users with similar music taste are manually adding to their libraries. A brand-new user to Last.fm will only have content that they search for or actively listen to. After a few days or even hours of listening to their favorite music, the user will receive randomly inserted music into their libraries. This is the result of the collaborative filtering technique looking up other users who listened to the same artists and finding out what other music they enjoyed as well. If it all goes according to plan, the new user will be satisfied with the recommended music and keep listening to

it. On the other hand, Pandora uses a content-filtering approach. They gather song attributes from their patented Music Genome Project database to find other songs and artists that overlap in their attributes. The database stores a mind-boggling 400 different attributes per song. This hodgepodge of different but similar music can then be refined by the user making a playlist. They simply “dislike” a song and the algorithm will deemphasize the songs attributes from the song selection process. The more a user dislikes, the more accurately the system can recommend the music they want to listen to. Likewise, if the user “likes” a song, then the algorithm looks for music with the same attributes that were liked. In Pandora’s view, all music can be boiled down to these 400 attributes.

With the current information glut online, it is no wonder that companies developed these systems. In theory, recommender systems are a win-win solution. The customer doesn’t have to sift through endless amounts of data nor do they have to suffer from information overload. The service provider, in turn, gets more profit. The main area of contention with recommender systems is that they might leak user information. Efforts to anonymize data in the past have met their fair share of criticisms by security and privacy advocates. Netflix was actually sued by a small group of individuals who learned that their anonymized data used in the Netflix Prize competition could be cross-referenced with free online resources to reveal their identities. Netflix went on to settle for an undisclosed amount. Other concerns in this space center around discrimination and possible misuse of collaborative filtering data. If this data isn’t anonymized, it can be linked to individual user accounts, meaning the company has private knowledge of the things you like. This, in turn, can fuel shady targeted marketing schemes.

The future of these systems points in a clear direction: every increasing level of personalization. Currently, recommender systems are isolated into the content world. Users like content, they can be recommended content and thus will want to use the service more. However, what if these systems could be generalized to areas that don’t focus on content? This technology combined with the internet of things, for example, could create the ultimate personalization climate for consumers. Recommender systems could connect to the user’s smart fridge and tap into a vast network of food preferences across all smart fridge owners. The smart fridge app could then recommend the user popular food items based on several attributes. Some of these apps could automatically place an order for them as well. Pandora uses 400 of

these attributed to songs; a similar system could be developed for food. You can probably already imagine some of these attributes. Carbohydrate content, sugar content, keto friendly, diabetic food, and several others could contribute to healthy customized diets.

Chapter 11: Robotics

Robotics is an interdisciplinary field that combines elements of engineering, science, physics, computer science, and electronics, just to name a few. Robotics concerns the design, implementation, programming, and maintenance of robots. A robot can be any physical system that is designed to perform an action or series of actions with varying degrees of autonomy. Robotics has a long history and has probably been in the human psyche ever since the first Greek stories of automatons were being told – though today’s robots are a little different from the ones you learned about in Chapter 2. Modern technology allows robots to be smarter, more lifelike, friendlier, and overall more useful than ever before. Many believe that humanity is headed for a new industrial revolution that will be spearheaded by the mass adoption of these systems in virtually every economic sector. Besides the obvious sort of industrial robots that you find in car manufacturing plants and refineries, there are also robotic systems being developed for retail, hospitality, transportation, telecommunication, medicine, and food services.

Robotics and artificial intelligence often get lumped together, but they are both distinct things. While some aspects of robotics lend themselves to use artificial intelligence techniques, the field as a whole is not preoccupied with simulating intelligence. Robotics can also be divided into several subfields that specialize in certain applications of robotic principles. On the industrial side of things, robots are used to handle tasks that are easy to delegate to a machine. Some of these tasks require high precision or repetitive motions that would be too difficult for a human to sustain for long periods of time. Others are designed for high throughputs like food packaging and labeling. These systems vary in how many behaviors they are capable of performing, their physical properties, and how they are programmed. Industrial robots come in different varieties, but one of the most common configurations is the “arm” type. These have similar designs but perform different tasks. They also have different degrees of freedom that allow them to move differently. Kinematics is very important for these types of robots because the mathematical equations determine how the joints on the arms move.

Interest in humanoid robots severely lacks with industrial robots regarding market penetration and usefulness. While research in androids or human-like

robotic systems remains high, the role of the humanoid robot remains a techno-curiosity more than a commercially viable business – though one day it is conceivable that they will be bought and sold as housekeepers like the robots in *The Jetsons*. There already exist housekeeping aid robots like the Roomba that has been available since 2002. The Roomba is manufactured by a robotics company fittingly named iRobot. Rodney Brooks, a professor in robotics at MIT, co-founded the company in 1990 with fellow classmates. Since then, Brooks has gone on to start another company called Rethink Robotics in 2009, most known for its creation of the Sawyer and Baxter collaborative robots. Baxter is fundamentally an industrial robot with industrial capabilities but with a bit of imagination. Unlike most industrial robots that usually resemble a mechanical arm, Baxter has two main limbs and an animated LCD face. It was designed to perform mundane tasks on an assembly line.

Baxter is a humanitarian take on the industrial robot. Instead of performing tasks rapidly and mechanically, Baxter uses sensors to become “aware” of its surroundings. The LCD screen will display a face according to the robot’s state. It can also respond to changes in its environment like ceasing operation if it drops a tool or piece and is unable to recover it. One of the most intriguing facets of the robot is that workers can programme it in the facility. Where other industrial robots require an engineer to be configured through control systems, Baxter is programmable by “hand”. An unskilled worker can move the robot’s hands to perform a certain task, and the computer will try its best to reproduce the movements. This makes Baxter accessible to everyone on the production line, not just educated technicians. Baxter has been lauded as being safer than traditional industrial systems that do not pay attention to other factors outside of their immediate programming. Despite this, the company went out of business in October 2018 due to low sales. Many old guard manufacturing companies saw these systems as experimental and the technology probably not there yet. The robots Sawyer and Baxter are still used in research today. Some universities use them to teach students in robotics courses. More recently, researchers hooked up electrodes between Baxter and a human operator to directly transmit brain signals. Sometime in the future, the interaction between human and robot may resemble something that is as intuitive as simply thinking.

Most industrial robots do not require to sense things or be particularly smart. They perform a set task and are usually left alone. In contrast, humanoid

robot research focuses on the ability of the robot to perceive the world around them and to interact with it. To be able to move a robot requires some propulsion system, like a series of actuators or electric motors. Another option is to use hydraulics or pneumatic systems. All robots need an energy source, either a battery or a direct connection to the current through a wall socket. To be able to see, a robot can be equipped with cameras, LIDAR, and various sensors. Computer vision is coupled with the cameras, and the sensors directly send environmental data like speed, position, balance, and so forth to the robot's control system. If machine learning is just a combination of statistics and programming, robots are a combination of a great many other things. In both cases, neither approach a general intelligence like the *Star Wars* droid 3-CPO. If that is the case, what is the current state of the art in robotics? Robotics research is split into many areas of focus, with some people working on sensors, robot dexterity, human-robot interaction, autonomous motion, and so on.

Humanoid robots are projected to disrupt areas of retail, hospitality, and food service. The chances are that you have already seen the automated menu systems for McDonald's and other fast food restaurants online. These systems are definitely being worked on, but may not see the implementation for various reasons. Some startups are even working on burger making systems that can perform virtually all of the duties of the average burger flipper at a fraction of the cost. The role of the humanoid robot in customer service is more obvious in places like Japan where robotics has entered the mainstream. Of all the industrialized nations today, Japan has the highest robot density by a large margin. Most of these robots are used in the automotive industry, but a number of them take on customer service roles. They can be found in select department stores greeting customers and in airports acting as luggage carriers. With an ever-contracting population and less unskilled workers willing to do these jobs, it makes sense that Japan has a high adoption rate. For many Japanese citizens, service robots serve to increase the quality of life and are a normal part of their daily lives.

However, there is still a long way to go for the average consumer in other countries to warm up to these robotic and automated systems. Some would even say that the need simply isn't there. In 2017, Japanese deaths outnumbered births by 1,000 to one. The population dropped by a massive 264,000 people in one year alone. Other industrialized nations with rising unemployment rates and a stable population perhaps don't need as many

robots. There are also other concerns besides just population and unemployment for why service robots are yet to be seen abroad. Overcoming the uncanny valley effect is no easy feat, especially with older generations. In Japan, people are used to them, so the effect holds less weight – though the same cannot be said for older generations in other countries. Young people today are more comfortable using self-checkout lines and automated systems, but the older generation overwhelmingly subscribes to the importance of face-to-face interaction with their community. A big argument for service robot adoption in Japan is their large aging population. End of life care is synonymous with assisted living, loss of personal autonomy, and embarrassing mishaps. These are areas that service robots can directly address. For one, an older patient is less likely to feel embarrassed if a robot is cleaning after their bathroom accidents. If there is a strong human-robotic interaction, the patient may feel that the robot is a part of themselves like we think our phones are. The result is a net gain in personal autonomy, all things considered.

With a shortage of geriatrics and personal care professionals plus a global aging population not just in Japan, service robots are projected to rise. Less and less people are signing up for these important ends of life careers. The pay for these jobs is low when compared to other medical professions that require a similar amount of training. If nobody else wants to take care of the elderly, who will? Not every family is willing or able to care for aging parents. Additionally, the risk of malpractice and abuse in assisted living institutions is high. Taking care of the elderly is not an easy job by any means. Combine this with a meager wage, and you have the basis for unfair treatment. Every aging person deserves to be treated with dignity and respect. Whether service robots can render this kind of treatment is up for debate. And it is a debate that will doubtlessly unfold within the current century.

Would older generations be keen on using a service robot for their daily needs? To some, a robot may seem intimidating or impersonal. After all, robots are not intelligent the way that we are. They may be able to handle basic conversational skills and respond to our commands, but that is a far cry from genuine intelligence. There is also the question of whether these systems should have a human-like appearance or to remain obviously robot looking. Humanoid robots are more likely to create an uncanny valley effect, but it is not exclusive to human-like design. Any robot that acts sufficiently “alive” can elicit the effect. Research into the area of the ameliorating, the

uncanny valley falls under human-robot interaction. We know from this research that certain robotic behaviors elicit certain human emotions like fear and uncertainty. For example, when a robotic agent gets too close or invades somebody's personal space, there is a fear response. If a robot is just lounging around with no purpose, the human perceives it as daunting and even useless. Interacting with these systems can result in unanticipated behavior in the human. Humans are likely to ascribe personality traits to a robot even when these aren't explicitly programmed. This factor of unpredictableness has lead researchers and designers to add emotive cues to their robots.

Where service robots suffer from the uncanny valley effect, utility robots do not. Virtually every major military across the globe is investing in robots to aid in the battlefield and in operations. A US-based company, Boston Dynamics, is working on robotic systems known for their mobility. The company made headlines with its design of a quadruped robot named BigDog for the Defense Advanced Research Projects Agency (DARPA). The robot has obvious military applications as a pack mule carrying ammunition and supplies for soldiers on patrol. The project was eventually scrapped because the system was deemed too loud to be used in live combat operations. These sorts of systems are mostly considered prototypes though. Other robots used by the US military include the Foster-Miller TALON family of remotely operated vehicles. They resemble little tanks or planet rovers with fully tracked movement. They are capable of using small arms fire to heavy machine guns completely remotely. Though TALON saw some deployment in Iraq and Afghanistan, its current use remains limited. Other fully weaponized systems are readily used like the General Atomics MQ-1 Predator drone and the MQ-9 Reaper. These belong to a class of aircraft called UAVs or uncrewed aerial vehicles. Predator and Reaper drones can either be controlled by an operator on the ground, or they can fly autonomously with the direction of onboard computers. These systems have faced increased scrutiny with high profile killings involving civilians and questionable rules of engagement. In 2011, under the Obama Administration, a 16-year-old American citizen of Yemeni descent was killed in a drone strike while eating at a restaurant in Yemen. It is unclear why the boy was targeted, but his father was a suspected al-Qaeda leader and was later killed in a similar strike.

The legitimacy of so-called "killer-robots" comes into question now and then

when a new system is announced. At the heart of the debate is whether a nation should possess the kind of power to kill remotely without trial and in the case of 16-year-old in a country that is not at war. Commentators are quick to point out that a machine with killing power is capable of being hacked. They appeal to the “Skynet” scenario where autonomous weapons gain sentience and begin to exterminate humans. Such scenarios are mostly far-fetched science fiction and are not helpful in the debate. While we know that some of these systems are vulnerable to getting hacked, we have yet to see a hacked autonomous system cause loss of life and limb. If there ever is a case where one does, it may cause enough trouble that governments are forced to ban them. Still, others question whether robots going to war is even a good idea. If state-of-the-art machine learning algorithms can barely drive autonomous cars with a certain degree of safety, how can we expect a robot to accurately wage war? War has several more dimensions of complexity than driving down the street does. Deploying actual robo-soldiers seems like an unlikely use of military spending with current technology – though you can rest assured that major governments are researching them.

A “killer robot” is a blanket term that is also unhelpful. Should a Predator drone be considered a robot any more than a tank with an autonomous turret is? Sentry guns have been in development for several years now and are regularly deployed into the battlefield. The most common of these are called close-in weapon systems (CIWS) and are used to protect battleships from missile and aircraft attacks. A CIWS is basically a large caliber cannon with a high rate of fire that is capable of shooting down missiles at 4,000 rounds per minute or higher. Unlike your grandfather’s anti-aircraft systems that may have required manual sighting, CIWS use radar guiding on a rotating mechanical platform to lock on targets autonomously. Not only that, but once “live”, the weapon system can engage without manual input after locking on. Any missile or aircraft that is detected by the radar system is automatically fired upon. In the case of the Phalanx CIWS used by the US Navy, a single 20mm Vulcan gatling cannon is used with computer-like accuracy. The problem of targeting is reduced to a simple algorithm. It doesn’t require machine learning or any other fancy artificial intelligence techniques. If the radar picks up an object that is approaching the ship, it needs first to make sure that its course is directed towards the ship. If the approach of the object is directly heading towards the ship, then the system must decide to fire. If the object is in between the minimum and maximum velocity range, then the

cannon fires. But if the object is too slow or too fast, the system does nothing. Both the minimum and maximum velocity range can be programmed by the operator.

CWIS, like the Phalanx, are considered autonomous weapons, yet they use very simple tech that has been around since the Gulf War. However, by today's standards, this is still a "dumb" cannon. The imperfect nature of the weapon has been demonstrated in a few disastrous training exercises. In one instance, the system successfully shot down a dummy drone target, but fuel and debris from the explosion damaged the ship and crew members. In another instance, the dummy drone was shot down, but the system reengaged it as it was in free fall, inadvertently sending rounds in the direction of an opposite ship and injuring crew members. This behavior should be expected from its simple rules of operation. A more sophisticated weapon system could probably have avoided these incidents. But as a last resort missile defense system, the Phalanx does a relatively good job against isolated missile strikes. Other autonomous, radar-guided systems are common throughout the world. Again, these don't need complicated machine learning schemes to identify targets and engage them. Israel's iron-dome system of missile defense is hailed by some commentators as the most advanced missile defense system in the world. These systems demonstrate some degree of intelligence, but really all they are doing is calculating velocities.

All things considered, the future of robotics is bright. Demand for robot systems will likely rise with uneven proportion to robotics professionals, creating well-paying jobs for those who are interested. There will be a glut of these jobs and a shortage of robotics designers and engineers. At the same time, lower-skilled jobs like robot technician and repairmen may emerge. Robotics will infiltrate other areas of tech like the cloud computing space and the internet of things. Robots connected to the cloud will lead to the creation of a cloud robotics marketplace where bidders can remotely program robots to perform certain tasks. Just like Alexa skills, these specific robot programs can be bought and sold in an open marketplace. Not only this, but robotics is taking on increasingly more important roles in business structure. A new executive position called chief robotics officer (CRO) will emerge for bleeding-edge organizations that make heavy use of robotic systems.

There are split views about just how far the field of robotics can go when it comes to automating human tasks. Whether or not the introduction of these technologies leads to technical unemployment and whether they should be

regulated is up for debate.

Chapter 12: The Internet of Things

Imagine a sea of connected devices that can all communicate with each other and transmit information. The internet of things is one of those buzzwords that gets thrown a lot, but that nevertheless has a relevant name. The internet is short for “network of intra-networks” or simply network of networks. A company or university has their own network of connected computers that may or may not be connected to the internet. This is called an intranet. When you connect those intranets through a common routing protocol as the internet does, you have digital communication on a massive scale. The internet of things then is a network of connected devices. Smaller and smaller network adapter cards allow devices to be smaller and fit virtually anywhere. To give an idea of how small these devices can get, a Bluetooth low energy weather beacon is about 30mm in diameter, 10mm thick, and weighs about 7 grams. Besides being minuscule, this device has a low energy usage and depending on configuration, can last for a couple of years before needing a battery replacement. Devices such as these have a small transmission range depending on what wireless technology they use to communicate. But since they are small, inexpensive, and massively produced, many devices can be laid down in an area of interest to form a relay network that can effectively increase the range of transmission.

The internet of connected devices is set to explode in the coming decades both as these devices become cheaper and as solar energy cells improve. The Internet Protocol (IP) under IPv4 namespace for IP addresses allow for a maximum of 4,294,967,269 addresses. IPv4 uses 32-bit addresses about 232 bits in total size. There are only so much different combinations of addresses that the protocol affords. However, the much newer and improved IPv6 protocol uses 128-bit addresses and as you can imagine offers orders of magnitude more addresses. In fact, there are about 10 to the 22^{nd} power addresses, which is a mind-boggling number. Part of the reasoning behind the implementation of IPv6 was that the world was running out of IP addresses to use for websites. Another reason was that the internet of things adds significantly more connected devices that all require unique IP addresses. With IPv6, the world is ready to stomach the burden of the internet of things. At a consumer level, the internet of things penetration is high in recent years.

The so-called “smart home” revolution has seen an introduction of smart devices, thermostats, garage openers, toasters, microwaves, televisions, and refrigerators that are all connected to the internet. These devices usually come equipped with a common interface or dashboard that is configurable through a smartphone. Why would anyone want to buy a common household appliance that is connected to the internet? It’s a good question, one that marketing teams all over the world had to grapple with when these products were being designed. The benefits of an internet connection are obvious for some appliances and less obvious for others. A smart fridge that comes equipped with a barcode scanner can keep an inventory of the household food supply, as well as note when products are about to expire. The complementary smartphone app can then maintain a database of all food purchases over the course of the year and offer basic analytics. More technically inclined users may seek a way to export their data for further analysis. As you can imagine, these products tend to lean on the expensive side and are marketing towards middle and upper-class households.

A more general definition called household automation refers to the devices, technologies, and control systems that aid in the home economy. At the most basic level, you have things like garage openers and clap-on lights. Somewhere in the middle, you have lighting systems, home theatre systems, and temperature control. At the high end, you get things that are only limited by the DIY spirit of the owner. A capable enough individual, say an engineer, can conceive of an automated pet feeder system that only requires food to be replenished every once a while rather than at every meal. Consumer irrigation systems can also come with a “smart” component, possibly a dashboard interface with options for setting up the watering frequency and so on. FarmBot is an automatic home gardening system that requires virtually zero effort to maintain. It mimics the architecture of a CNC milling machine but is instead equipped with trowels and seed dibblers. While the basic kit costs upwards of \$3,500, the company maintains that the system is 100% open-source. This falls in line with the DIY ethic behind many of the internet of things products. Sales of home security systems have also been rising in recent years. Many technology companies are taking full advantage of the internet of things revolution in security spaces and consumers have multiple brands to choose from. Smart assistants like Amazon Echo and Alexa have also gained popularity. They aid in the home automation process by listening to user commands and executing some function. This may be to surf the web,

create a shopping list, play music, and a myriad of other things. The user can link up “Alexa skills” that are the programmable scripts that the devices run after each command. These skills are readily searchable online, and anyone with the programming knows how to publish them.

Home automation is especially relevant to the aging population as the needs of the elderly are myriad when living at home. For some, systems may delay the need to be admitted to a healthcare facility. However, with current technology there is only so much these systems can achieve. Stairlifts for those who use wheelchairs have been around for years. Home automation then isn’t a new thing, but it certainly has gained renewed interest with the internet of things technology. An application of the internet of things principles to home automation for the elderly has many benefits. One of the most important features that many elderly homeowners need is an alert system that can call emergency services if they are incapacitated. Non-emergency alerts like reminders to take medication and set up doctor’s appointments are also helpful. The true use of the internet of things methodologies would include a smart device that is worn on the wrist or chest that sends information to the patient’s doctor about heart rate and blood pressure. Recent advances in smart fabric technology allow this functionality to be embedded into the very clothes that the patients are wearing. In the future, these systems will be able to integrate with robotics for additional assistance. Domestic robots will prepare meals, clean after the patient, and assist with doing chores.

Another novel family of the internet of things applications belongs to the industrial sector. Smart sensors can relay information about weather conditions, equipment status, and logistics. RFID technology can be used to track stock keeping units (SKU) that move from warehouse to warehouse. The so-called “industrial internet” combines networked devices, big data analytics, and real-time updates. These always-on metrics are relatively cheap for large companies like General Electric to implement and yet they provide a wealth of information. Areas where the industrial or manufacturing process cause inefficiencies can easily be spotted by their operators. Once a bottleneck is identified, the necessary changes to correct them are given center stage. Increasing the number of connected devices lowers sunk costs in the form of productivity losses by optimizing the entire process. When you have all the data in the world, you have options. Though General Electric remains an industrial internet powerhouse, the concept has seen slow

adoption globally.

The internet of things devices have also been proposed for the creation of “smart grids” for common utilities. The idea is to use monitoring devices to gauge the level of electricity needs so that the grid can move resources to areas that need it most, meanwhile boosting efficiency. This enables two-way communication between the utility provider and the consumer. Smart grids ultimately lead to lower electricity rates for the consumer and general availability for all. A hidden benefit of transforming power grids into smart ones is that old equipment is replaced with the new. It is no secret that US infrastructures are slowly crumbling away. Installing new networked devices gives policymakers the excuse for finally getting rid of troublesome parts. A newer grid is a safer, more reliable grid. A smart grid is an informationally rich way to do utilities. The customer will have access to smart meters that can be accessed through their smartphones whenever they wish. There they can see their monthly bill, usage rates, and opportunities to save. A smart grid prioritizes resources for high-demand peaks, but at a greater cost. This allows thrifty customers to save their most demanding electricity usage for low-demand periods, saving them money in the process. Another benefit to smart grids is higher availability for vehicle charging stations. As it stands, the US power grid is not ready for the switch to a mostly electric car society. Instead, electric drivetrains are only practical in certain high-density areas where charging stations are readily available.

Since the internet of things technologies generates massive amounts of data, existing machine learning systems will only get smarter. If a company is already using sensors and other connected devices to harvest data, you can be sure that they are using a data lake or data warehouse solution to store it. Real-time analytics requires high bandwidth data ingestion engines to parse the information as it is being relayed to servers. Big data is only projected to grow alongside with the adoption of these networked devices. The internet as it stands right now generates a massive amount of data. Social media apps like Snapchat and Instagram upload thousands of user-generated photos and videos every minute. Facebook, LinkedIn, Quora and others generate data through user posts, as well as in-house and third-party analytics for each user. Virtually any high-traffic website is using third-party tracking and data harvesting add-ons. These record things like mouse movements, keystrokes, and clicks on advertising banners. Basically, any action that a human agent does on a logged in account can generate data for various companies.

If you add to this the internet of things, which generate different types of data formats depending on the wireless technologies, the global data glut shoots up. For example, the World Wide Web uses various data formats that are well known to data scientists. These include JSON, CSV (comma separated values), and XML. Many devices that are connected through web protocols like HTTP communicate with these types of columnar data. The number of connected devices is expected to reach 31 billion in 2020. That is almost 30 times the number of people who are currently online. The amount of data these connected denizens generate was somewhere around 2.5 quintillion bytes a day in 2017. We can expect the field of machine learning to grow, with new techniques and algorithms added to the already diverse repertoire of narrow artificial intelligence. Because connected devices can be virtually anything, it is difficult to say in what ways they will be implemented. We do, however, have a general sense in where things are going. The most important principles of the internet of things technologies are sensing, communication, and relaying. Therefore, the internet of things will be used to create large-scale communications systems of small devices.

A smart city is a type of urban metropolitan area that uses the internet of things connectivity for analytics, data optimized infrastructure and utility delivery, and transportation networks. A data orientated city is one that utilizes smart grid technologies to boost energy efficacy, has an open attitude towards civic participation, and has optimized fleets of public transportation. Citizens can download an application on their phone that gives them their own personal dashboard into the city in real time. They can read metrics, look up transportation schedules, traffic alerts, and several other functions. Technologies will allow them to find empty parking spots, report potholes, and gauge the human density of their favorite hangout spots. Most of this functionality will stem from connected sensor networks. Vehicular networks, including in-vehicle and vehicle to vehicle communication, will make driving easier, safer, and will pave the way for fleets of autonomous driving vehicles. A modern vehicle of the year is equipped with countless ECU (electronic control units) that each communicates with each other in an in-vehicle network. These are used for on-demand vehicle diagnostics, conformance to emissions control standards, and even braking systems. Vehicles are increasingly being computerized to the point where models that use drive-by-wire are common. These models use a type of electronic steering that obviates the need to have a mechanical steering column. Vehicle

to vehicle networks allows for efficient throughputs in busy roads. They can even be configured to detect accidents and avoid them in real time. If cars can communicate their positions, speed, and direction, then deterring an accident is a simple means of the car braking or steering by itself when the driver isn't paying attention.

Just like the other applications of these technologies discussed in previous chapters, the role of regulation and legal frameworks cannot be stressed enough. Even if the technology is there, it doesn't mean that there exists a solid policy for implementing them. Driverless cars are a regulatory nightmare, so are vehicular networks that alter the behavior of crewed vehicles. Something that is brought up time and time again with connected devices is the availability of secure communication. For many in the security industry, connecting vital infrastructure and private networks poses a security risk that outweighs the potential benefits. Since this is a relatively new space with different wireless technologies, potential attack vectors are everywhere. It is difficult for the internet of things to maintain a strong security posture because there is greater overlap between physical and wireless intrusion. Besides wirelessly attacking a sensor or vehicular network, a potential attacker can also go to the physical site where the sensor is located. In industry terms, these are called physical layer and networking layer attacks respectively. Since these networks use many individual nodes, the attack surface is only increased. The sheer number of consumers the internet of things devices connect to the internet poses the threat of virus proliferation. The threat of massive scale denial of service attacks is a commonly cited fear by security researchers. Imagine a zombie network of millions of domestic appliances like toasters, refrigerators, speakers, and thermostats that can use web protocols to inundate legitimate web services with communication requests.

Only time will tell if the internet of things security posture can be hardened to the point that regulators are more willing to accept them. Continued research in the safety of vehicular networks will have to improve for municipalities to give the go-ahead on roadside vehicular sensors. The same can be said for smart grids. All across the globe, experiments in smart cities are being conducted as you read this. A simple Google search can point you in the right direction towards such a city near you. Whatever the case, you can expect to hear more about the internet of things' security, implementation, vehicular networks and regulation in the coming decades. Any new technology is

usually slow to be adopted, but when it is, it has the potential to revolutionize society.

Chapter 13: Why AI is the New Business Degree

Once upon a time, the most popular college degree was a business degree. If a student didn't know what field of study they wanted to go in but still wanted a decent job, the business degree was the way to go. Now, there is an overabundance of students going for their MBAs but not enough going after degrees in artificial intelligence. While the business world is facing a saturation of business knowledge, artificial intelligence is facing a shortage of AI know how. There has never been a time in history where getting a computer science degree focused on AI methods has been more lucrative. Many top tier schools are now offering new degrees that focus on AI and the gamut of robotics rather than the general computing knowledge taught in traditional computer science programs. This means that there is market demand, as well as plenty of candidates interested in the degrees.

It is no wonder why AI is quickly changing the face of business in multiple sectors. Technology related jobs are increasingly asking that their candidates have a solid grasp of machine learning methods along with their other expected duties. Virtually every major AI player has already open sourced some of its machine learning libraries so that everyone from startups to large corporations has access to create AI programs with little upfront costs. Google released TensorFlow in 2015, and it has since become one of the most popular repos on GitHub, the definite open source authority today. Another open source library called PyTorch is used extensively at Facebook and Uber. While smaller players in the AI space benefit from hiring PhDs in the field, they already have many of the core infrastructures for creating neural networks provided by these free tools. The availability of these tools is further coupled with easy access to computational resources provided by cloud providers like Amazon Web Services and Microsoft Azure. A company no longer has to invest in a high-cost machine learning cluster of GPUs when they can simply rent as much processing power they need from the cloud. While these services are more expensive in the long run, they still make it easier for a company than building machine learning infrastructure outright. AI tools allow a business to engage in machine learning that before may not have found a use for it. A relatively new trend is to use automated chatbot software on their customer-facing web pages for quick and easy information

retrieval. A customer can ask for rates, available products, and other information by simply typing a few text commands. While chatbots are a bit behind in terms of realizing their full potential, many consumers are becoming more familiar with them. Some companies even hire human agents to fill the chatbot role until the technology exists to allow chatbots to perform these duties on their own. Another common application of AI technology is to predict customer behavior. The focus on analytics is at an all-time high for all types of companies, not just retail. Coupled with social media marketing and e-commerce, analytics drives better customer decisions for the company over the long run. There has also been a rise in purely “AI” focused companies that market a single product or line of products that have AI functionality. One of these is called Grammarly, an online service that uses AI methods to streamline the writing process. It provides editing for simple mistakes, writing pitfalls, and suggesting things the user can say to sound more professional. It is essentially a writing tutor that the user can take with them wherever they go. Yet another company called Stick-Fix uses AI to recommend clothing options to its customers based on a series of preferences. A user designates their price range, enters their measurements, and chooses a style they are going for, and the service sends them a box of outfits to try on. Large companies may use AI to automate their systems. This is especially true in manufacturing and human labor-intensive jobs – though many of these jobs will take years before being fully automated. The recycling plant example given earlier in this point remains to be solved on a massive scale. Most automation systems first begin by aiding line workers rather than replacing them outright. While low-skilled workers are at high risk of losing their jobs to automation, the shift will not occur overnight. The most likely scenario is that when robotics is first introduced into new industries, they will work alongside human employees. This is already true in the car manufacturing business. In 2018, Tesla Motors was highly scrutinized by business leaders in their attempts to automate large portions of their Model 3 production. The result was an overestimation of automation capabilities and an underestimation of human worker capabilities. The company undershot how many Model 3s they could output per month using their heavily automated plant. This move was criticized by industry veterans, some of them calling it a “rookie mistake”. The state of modern robotics is still behind the power of the mind and especially human dexterity. Mimicking the same micro-muscular movements used for manipulating tools in the human hand is

a non-trivial problem to solve with robots. It will likely take decades before a machine matches the human level dexterity of a line worker with years of experience in their craft.

However, it isn't just line workers who are in danger of losing their jobs. Advances in AI, especially in speech recognition and language processing, threatens to displace the huge call center industry with automated systems – though that is also likely to be decades in development. Retail is seeing a transformation with automated systems on top of the already high online sales numbers. Fewer people are going out to shop. At least in places like Japan, retail facing robots are actively being invested in. Even those who enjoy a comfy white-collar job have reasons to up their education and technical skills. Automated systems for payroll, accounting, and balancing the books are under active development. In the future, even programmers will not be safe from the AI deluge. Systems are being trained to perform the most mundane programming tasks that are often handed to lower-skilled coders. This includes software testing and looking for bugs – though, with all things under threat of automation, these systems will likely be deployed to work side by side with the programmer rather than replacing them altogether.

Driverless technology is also on the rise. Even as you read this, you can be sure that several companies across the world currently have autonomous driving systems on the road somewhere, gathering ever more data to strengthen their algorithms. What will likely happen with driverless technology is that the ability will come first and policy second. Just because driverless cars are proven safe, it doesn't mean that they will automatically be introduced. There are far too many gray areas to see the mass adoption of these technologies soon. In 2018, the first-ever driverless car fatality was recorded in Tempe, Arizona, by a car owned by Uber. You can be sure that more such fatalities will follow until the technology is perfected and here lies the difficulty in policymaking. There has never been a time in history when machine ethics have been used to determine laws. Even if driverless cars kill several hundred people a year (and they probably will), lawmakers and insurance companies have to decide if it is preferable to killing several thousand. There is an increasing need for stakeholders to have this ethics conversation, as well as for common citizens to be informed regarding policy and current advances. The major deterrent to driverless vehicles hitting the mainstream will be the law, not the capabilities of technology.

The industry is currently hungry for people skilled in the top AI libraries. The

highest paying positions are looking for PhDs and Master graduates, but a good portion of them are looking for anyone at all who are competent programmers. The need spans other industries besides pure software engineer as well. Mechanical and electrical engineers with knowledge of machine learning techniques will be in high demand for the coming years. Add to this a cursory knowledge in the internet of things, and you have a highly desirable candidate.

Chapter 14: AI FAQ

Question: Does machine learning have limits?

Answer:

Machine learning certainly has its limits. Even when advanced techniques like deep learning are used, these systems are limited by their data and feature selection. Deep learning is a little ahead of that curve because features can be selected automatically. However, machine learning can only be applied to the five general problems discussed in Chapter 6. If a problem lies outside of that problem space, then current machine learning will never be adequate for answering it. Simply throwing more data will be useless if machine learning doesn't align with the problem. Fortunately, or unfortunately, we haven't had machine learning robot legislators anytime soon.

Question: How will AI be used in the military?

Answer:

The US Pentagon has increasingly shown interest in the term “algorithmic warfare”. That is the application of machine learning methods to the battlefield. It will most likely be used for target selection and prediction of enemy movements. Autonomous systems with “kill decision” capabilities will no doubt be developed from these general applications – though as previously mentioned, there is a huge backlash against the militarization of AI by the research community. These weapons may or may not gain unconventional status by the international community somewhere down the line. Even if they do, major powers will still pursue their creation.

Another possible application of AI will be towards cyberwarfare, the infiltration of nations through subterfuge, and the targeting of autonomous weapon systems of enemy states. AI algorithms will only add potency to cyberattacks by state actors, with AI systems possible taking over enemy drones and uncrewed gun platforms.

Finally, there is the looming threat of facial recognition used to identify targets. It would not be out of the ordinary to suppose drone attacks in the future will be targeted at domestic actors. The drone strike in Yemen first sent shockwaves through the international community because it affirmed its right to strike at targets over foreign soil. We have yet to see such a strike inside the country.

Question: Should I be worried about losing my job?

Answer:

The short answer is no. Unless you belong to a certain market vertical where automation has already phased out many jobs, you are probably safe for the next ten years or so. If you are in a profession that is actively an area of AI research like the driving industry, you are still protected by the buffer zone called regulation. If you take your job seriously and you belong to these industries, it will be worth staying up-to-date with current technological advances and especially regulations. The first introduction of automation technologies always aids the worker before phasing them out completely. Truck drivers already benefit from cruise control, for example. Driverless trucks may still need operators on board to oversee functionality and extreme edge cases the computer cannot resolve. What happens if a crazy person jumps in front of the truck to mess with it? They are more likely to do this if the truck is completely uncrewed. The same goes for potential hacking scenarios. Thankfully for you, driverless cars are a regulatory nightmare and probably won't see a mass adoption for a few decades.

Question: I've heard that AI can drastically change the world, bring about world peace and even end poverty. Is there any truth to this?

Answer:

It really depends on whom you ask, as nobody understands what AI will look like in 30 or 50 years down the line. Artificial general intelligence that converges with humanities goals will be able to perform many tasks better than we do. It can then work on these problems without needing to eat or sleep, bringing about advances at an exponential rate. Such systems could probably solve the most pressing problems today, like climate change and food logistics. Some believe they will bring about a new age of abundance where people are not required to work anymore. The wealth created by this intelligence will fuel a universal basic income. Those are the most optimistic projections. A more level head explanation may be a world where general AI is limited to select corporations and governments, further contributing to wealth inequality. Still, another practical explanation is that general AI is a fad idea that will never see reality.

Question: How worried should I be about AI in general?

Answer:

For the average person, there is little to worry about. Most applications of machine learning are completely benign. There are some privacy and ethical

concerns, but then again, most people will not be affected by them. Even if the military adopts these weapon systems, there is an asymptotic risk that they will be used against you. Which is to say, virtually zero. Losing a job may be a concern, in which case you have a long head start to think about what other industry you can switch into. Remember for every terrible application of the technology there is a good potential benefit to humankind. Imagine a world where graphics content is removed from family websites like Facebook. Imagine a world where driving is as safe as flying a plane, where young drivers are not at an increased risk for fatalities. Imagine a world where market inefficiencies are smoothed out, leading to an increase in the economy. Imagine a world where cancer is no longer a major killer due to early detection and advanced drug discovery. Imagine a world where the elderly live longer and are cared for with the dignity and respect that they deserve, no longer having to worry about abusive caretakers or the embarrassment of loss of autonomy.

Question: If general AI is discovered and it does go bad, what can humans do?

Answer:

Like with most intractable scenarios, prevention is the cure. Trust that companies are doing the right thing to ensure that their systems do not run off course. In the future, AI legislation will be a common topic of debate. If you are not currently a voter, you may reconsider when the fate of humanity is at stake. There is a striking hypothesis about the creation of general AI that posits that the sooner it is discovered, the less of an existential threat it will be for humanity. This is because the longer it takes, the better technology and systems it will have at its disposal. Imagine if nanotechnology 3D printing machines are in active use when the general AI comes online. If it were to go rogue, it could very quickly begin transforming the face of the Earth into a massive computer.

Conclusion


Thank you for making it through to the end of *Artificial Intelligence: What You Need to Know About Machine Learning, Robotics, Deep Learning, Recommender Systems, The Internet of Things, Neural Networks, Reinforcement Learning, and Our Future*. It should have given you a better understanding of AI. The hope is that you can understand internet headlines that talk about AI without feeling completely clueless. The knowledge contained in this book should have been enough for that end, and to encourage making your own conclusions.

The field of AI is vast and ever-changing. Many methods of AI like genetic algorithms, statistical inference, and stochastic processes were not covered in this book, but there are countless resources out there that do. Additionally, you can find a complete history of AI on the Wikipedia page with current developments as well. If you wish to pursue artificial intelligence as a course of study or even as a career, be prepared for lots of math and computer programming.

Finally, if you found this book useful in any way, a review on Amazon is always appreciated!

Part 2: Internet of Things

What You Need to Know About IoT, Big Data, Predictive Analytics, Artificial Intelligence, Machine Learning, Cybersecurity, Business Intelligence, Augmented Reality and Our Future



Introduction

You were just woken up in the middle of the night by smart lightbulbs in your house blasting at full power for no reason. Your bleary-eyed investigation shows they tried to download a firmware update and failed. At that moment, Alexa starts quietly whispering sweet nonsense to herself in the corner and Roomba starts slamming into the nearest wall. What do you do? Is your house haunted or have the machines finally started an uprising? Neither – it's just another day in the IoT wonderland.

The following book reveals the concepts and methods powering perhaps the most ambitious technological concept of the twenty-first century – the Internet of Things (IoT) – and parades all the ridiculously named gadgets techies imagined to saturate the market before the competition. Mystical, cheap and scalable, the idea of IoT attracts creative grifters of all shapes and sizes to try their luck in pushing yet another completely unnecessary gadget to the market in hopes of fleecing gullible buyers. What you're about to read contains all the juiciest examples of IoT technology, including:

- Smart sprinklers that can be turned on and off from halfway across the world
- A smart toilet with ambient lighting and speakers for total immersion
- Smart scented candles with the scent of money set on fire
- A smart fishing rod for gathering stats on the spot
- A smart air purifier that moves around the house
- A smart water faucet with LED lighting
- A smart menstrual cup
- A smart block of wood
- Smart bumblebees
- A smart padlock
- Smart farms

IoT shows amazing potential in medicine, where it can relieve doctors and nurses of daily drudgery related to managing chronic diseases, such as diabetes. In the meantime, satire and wishful thinking abound in IoT, presenting us with a glorious reality full of humor and head-scratching. *What were they thinking?* Well, let's find out.

Chapter 1 – Origins of IoT

You know that old song with the lyrics, “Foot bone connected to the heel bone; Heel bone connected to the ankle bone” and so on? If you imagine a vast, digital body spanning the entire world whose parts are connected just the same way as described in the song, except they’re made of information and tiny gadgets, you’ll come very close to the idea of IoT. Packets of data traveling back and forth in the global IoT body would then represent the nervous activity in a living body, where cells communicate with one another to coordinate and fulfill some greater purpose for the benefit of the entire organism. The definition of **IoT** would thus go: a series of devices with some ulterior purpose that has been given internet connectivity.

It’s hard to say who, if anyone, conceived the notion of IoT but we can guess busy scientists looking to shave a fraction of a second off their interactions with the real world were the first to embrace the idea of interconnected gadgets and bring it to fruition. Because they didn’t care about flashiness, their gadgets were crude and their protocols efficient, which minimized **attack surface**, joint weakness of a network that directly correlates with complexity. There is a way to deploy IoT in the real world securely, but it has to be done by professionals who know the risks and benefits.

During the 1970s, scholars of the Massachusetts Institute of Technology (MIT) were enjoying cold fizzy drinks from the campus’ Coca-Cola vending machine at a discount price. The problem was that, as the campus grew, the drinks would be snatched up almost immediately by passersby at the expense of those on the outskirts of the campus, who had to walk ten or fifteen minutes just to find there’s nothing in the vending machine, or worse yet, that it’s just been refilled, and the sodas are still warm. So, there was an actual problem in a tightly-knit community that led to frustration and loss of productivity. As we’re about to see, IoT *can help* in situations like these.

The vending machine received micro-switches in each of the six columns to keep tabs on when each bottle was placed inside and whether it was sufficiently chilled; after three hours, the central processor would mark the bottle as “cold” in the companion program. The vending machine was given its own user account in the internal campus network, allowing anyone to ping for user “coke” to check out the bottle status. Anyone hooked up to the internet who could access the campus network was capable of checking on whether bottles were chilled, though there wasn’t much use in the function if

you were halfway across the globe^[1].

Note the organic process of how IoT became integrated with existing technology – members of a tightly-knit community were experiencing discomfort and loss of productivity due to outdated technology that provided insufficient information. By integrating small and highly specific IoT capability in the existing technology infrastructure, discomfort was averted, and productivity losses were minimized. There are plenty of ways to sabotage this particular IoT implementation, but any such saboteur would need physical access to the campus, in which case he or she will be easily caught or identified. This is *not* how IoT will work for the general public. Instead of having specific features requested by customers, IoT devices will have a plethora of gimmicks that will open workplaces and homes across the world to relentless hacking attacks.

No single entity decided to create IoT; it's actually a spontaneously emerging network of loosely allied devices. Software and hardware industries are *aching* for a set of standards, and IoT seems to be the closest we'll get to a global interconnectedness standard. So, shower curtains from China, wool socks from Italy and coffee mugs from Argentina can all be given internet connectivity to turn them into IoT devices that can communicate, but the question is – why?

IoT actually allows companies to offset a part of their production cost by gathering and selling data of customers, hiding inflation in the process. The intrusion of privacy is still there, but it's much easier to ignore it when you appear to have saved 20% on the price of a shower curtain or a coffee mug. Sure, you'll be presented with a privacy policy or terms of use where the fine print states your use of the product will be tracked, but who reads those? When was the last time you read a 'terms of service', let alone the fine print in one? When a company sees that customers don't care about their privacy and starts blatantly spying through IoT functionality, then all companies have to start doing it or risk falling by the wayside.

On the practical side of things, IoT shower curtains can measure humidity in the bathroom and automatically open IoT windows to let some water vapor out when you're done showering. You did buy IoT windows, didn't you? IoT wool socks can measure circulation in your feet and alert when you should stretch or go for a walk, and an IoT coffee mug can display the time on its surface or just communicate with your IoT coffee table to warn you through an app that your coffee is getting cold. You did buy an IoT coffee table,

right?

See how it works? Each IoT product provides a crumb of utility that only comes true when you buy the missing ten items that confer additional functions to the entire set. Oh, and buying another ten items unlocks this capability and another 100 items that one and so on. By acting upon the hoarding instinct inherent in all humans, companies producing IoT gadgets aim to make a fortune by peppering our living spaces with seemingly useful gadgets that do have a marginal advantage over non-IoT items in the same category but are otherwise just the same, except pricier and completely insecure.

Chapter 2 – IoT Security

In any network, the main principle regarding security is that the entire network is only as secure as its weakest link. How much security do you think an IoT shower curtain would have? Precisely zero because the manufacturer, most likely a Shenzhen factory, will be looking to minimize production costs and offset any liabilities onto the next link in the chain, such as IoT windows. Well then, how much security would IoT windows have? None whatsoever because the manufacturer of those would again be following the same logic of making the biggest profits.

Once a hacker hijacks any IoT device in a household, he or she will gain possession of them all in a cascading manner, turning the entire IoT network against the owner. Foot bone connected to the heel bone; Heel bone connected to the ankle bone... In one instance, a casino got its entire network compromised thanks to an IoT thermometer in a fish tank in the lobby^[2]. The entire database of players was hacked into, copied and pulled out through that same thermometer and nobody was the wiser. It can take years before anyone notices these breaches of security, and even then, people in charge are likely just to shrug it off.

There is simply no set of IoT safety standards like there is with food, cars or bikes; anyone can make IoT devices and market them globally without any liability. There is no insurance against hacking attacks either, making the entire IoT field a haphazard endeavor, which is great for entrepreneurs that have nothing to lose but pretty rotten for regular people and businesses who are hyped for technology that end up getting burned. That doesn't mean IoT is useless, but simply that it has to be deployed in an environment that is already secure and with well-known, trusted users, just like we saw with MIT scholars and their soda machine. What happens when IoT is deployed insecurely worldwide? Hacking attacks, the scale of which dwarfs everything we've seen so far.

In 2016, a massive wave of internet signals smashed at the shores of US consumer devices and caused huge congestion in traffic. It came from IoT devices carelessly left online for anyone to hack into and take over. CloudFlare, the intermediary company that analyzes internet traffic and mitigates cyberattacks, studied this particular bout of distributed-denial-of-service attacks and found they were mostly coming from Vietnam and Ukraine^[3] but were otherwise carefully orchestrated across thousands of

different IP addresses. At times, the volume of traffic went over 1 million requests per second and consisted of 52,467 unique IP addresses. Analysis of attacker traffic showed that Vietnamese devices were most likely CCTV cameras due to ports they had opened.

IBM's Chief Technology Officer (CTO), Bruce Schneier, warned in 2017 that governments have to take IoT security seriously and step up their game before the damage is done. He said, "We're building a robot the size of the world, and most people don't even realize it" in a keynote address at the SecTor security conference in November 2017^[4]. What used to be cybersecurity now has to rapidly evolve to become "everything-security", implying there's no stopping the building of IoT, but we can at least minimize the vulnerabilities.

In January 2019, the Japanese government announced an IoT security project^[5] during which sanctioned hackers will be scanning IoT networks and trying to breach routers and webcams on vulnerable networks in preparation for the 2020 Olympic Games. Hackers will be using what's known as a **dictionary attack**, meaning they will have a compiled list of all the most commonly used username/password combinations, such as "admin/admin" or "admin/blank". The list will presumably be forwarded to Japanese ISPs who will then alert the owners of those devices to change usernames and passwords. See any weakness in this plan? If a black hat hacker were to get a peek at that list, it would make their work a whole lot easier.

Why such fear? In 2014, the Winter Olympic Games held in Sochi, Russia, were protected by the heavily entrenched Russian army with over 40,000 law enforcement officers. Security checkpoints were set all over the place with X-rays and metal detectors while aerial patrols circled the skies and gunboats patrolled the sea. What about cybersecurity? Internet traffic was thoroughly analyzed, but Sochi was still besieged by hackers that set up plenty of traps for naive tourists who just wanted to get drunk and roll in the snow. For example, after arriving at the airport, a tourist gets a notification that there's free Wi-Fi access as long as he or she downloads and installs a special app.

The trick is that the app is actually malware and captures passwords and banking information. Even when the tourist goes back home, he or she will often keep the app, which will continue to leech private information. Even if someone were to notice something fishy and figure out the app did it, can you imagine calling the police because of a malicious app? You'd probably get arrested and fined for being a nuisance. In this way, hackers use a low-risk,

high-reward strategy that exploits laziness of a general smartphone user.

Anyway, that Russian athletes were doped up during Sochi Olympics came out later that year, with the governing body banning all Russian-affiliated athletes from the 2018 Olympics, which is when someone, supposedly Russian state-sponsored hackers, tried to hack the South Korean Olympic games. Dubbed “Olympic Destroyer”^[6], this particular strain of malware was well prepared by someone who had inside knowledge of systems in place. Olympic Destroyer would nestle on an infected machine, steal passwords in an attempt to infect the entire network and then deliver coup de grace by completely wiping out everything from the machine, including any trace of infection. This led to some disruption to the opening ceremony and the Wi-Fi network in use, but otherwise, everything was smooth sailing.

Cybersecurity researchers later claimed that Russian hackers made Olympic Destroyer, but then another group of researchers said it was Chinese hackers. Well, which was it? Nobody can tell. All hacking attacks leave behind traces of information, but it’s impossible to know if the hackers were simply sloppy or played mind games with researchers. That’s the scariest part of IoT – the fact you could be sitting in your cozy Wyoming home playing “Fortnite” while Danish and Estonian hackers try to disrupt each other through your network, using your devices and spending your power to mine cryptocurrencies or make DDoS attacks. Unless you have the cybersecurity expertise that can at least match hackers, you’d be none the wiser, and you’d be footing the bill. At least California is doing something to stop an IoT disaster.

In September 2018, California governor Jerry Brown signed SB 327^[7], a cybersecurity bill meant to tighten up IoT security, slated to go into effect January 2020. Current cybersecurity law requires that a California business undertake “reasonable security procedures” to maintain the privacy and security of its customers; SB 327 aims to expand that to “reasonable security feature or features that are appropriate to the nature and function of the device”. Whoop-de-do. Why are legislators so afraid of locking down digital technologies? The answer to that lies in China.

The Chinese government has quite an interesting mentality – economic victory at any cost. To match that the US government simply has to abrogate the constitutional rights of its citizenry, at least when it comes to making free market choices. By giving domestic companies plenty of legal and economic leeway, the US government foisted them as de facto arbiters of right and

wrong in the country. This is why the 2008 credit crash in the US resulted in taxpayers bailing out banks that went all in and lost horribly; without a bailout, the Chinese would have swooped in, bought them off and then it would have been game over. So, anything China does, the US has to do to an extent or risk falling behind.

When applied to IoT, what this means is that regular citizenry will have their privacy invaded the same way Facebook and other such companies already do in order to generate value and stay competitive in the global economy. Sure, there will still be laws like SB 327, but they'll intentionally have loopholes so US companies can compete and be applied only when the general public gets so irate that it needs a scapegoat. If you want some peace and privacy, you'll have to hack your way to it.

Chapter 3 – Ethical Hacking

In the context of IoT, ethical means “distinguishing between good and evil” and hacking means “unorthodox use of a system or tool for a palpable advantage.” Without going into philosophy, good and evil refer to having a goal; good is whatever brings you closer to that goal and evil whatever makes you stray away from it. So, if your goal is to have privacy, then ethical hacking helps you achieve privacy through unorthodox use of systems or tools. Sound good?

One example of ethical hacking is overclocking graphics cards. In essence, manufacturers of graphics cards for desktop computers typically cap their strength anywhere between 80-95% of their full potential. By tinkering with graphics cards, it’s possible to remove the cap on the internal clock they’re using (hence *overclocking*) and unlock the performance that’s already there but hidden away from the user. Thus, ethical hacking gets you what you paid for, but the company doesn’t want to provide for whatever silly reason.

Keep in mind that the US government considers hacking a big no-no, but the legislators are mainly older people out of touch with technology who consider the internet “a series of tubes”. As long as the ethical hacking you do isn’t a nuisance and doesn’t do harm or economic damage, you’re pretty much under their radar. This applies to law enforcement as well, which is typically so overwhelmed with traditional crime that it has no time to deal with eccentrics tinkering with toys in their garage; again, unless you’re being a nuisance or doing harm or damage. Don’t attract any undue attention to yourself and keep working on your IoT customizations, which is all ethical hacking is.

Conversely, all hackers prosecuted in the US to date have had the charge of wire fraud levied against them. The definition of wire fraud is so mind-bogglingly broad that it includes intentional misrepresentation of fact to achieve deception through electronic means of communication. Essentially, girls posting their selfies with camera filters and puppy ears are committing wire fraud because they’re misrepresenting their faces to deceive observers into thinking they’re cuter than they are. Thus, if law enforcement wants to make an example out of you, they can find so many ways to do it.

Companies churning out IoT products are in a similar bind – they have deadlines to meet, half-baked products to push out and lawsuits to fend off. Everything they do is highly optimized to deliver maximum revenue. They

have no time or resources to deal with each ethical hacker unless he or she is, you guessed it, being a nuisance or doing harm or economic damage. This leaves you plenty of space to actually get some of those God-given constitutional rights without interfering with companies on their revenue-gathering rampage.

The exact specifics of ethical hacking are a bit trickier. IoT technology advances so quickly that it's truly a thankless task to write any kind of tutorial on it, especially one that's meant to stand the test of time. However, technology is typically improved incrementally, meaning some underlying principles are likely to apply to several generations of IoT devices.

Ethical hacking is about taking things apart and watching them tick. So, when you get your hands on an IoT device that you won't feel sorry to see die, take it apart in controlled circumstances and watch how it ticks. This will show you how companies assemble their products and also underline how shoddily they're built – what with third world factories hastily assembling them for pennies. So, if you can take your time to understand any given IoT device and improve upon it, you've become an ethical hacker. Preferably, also disconnect it from the internet and don't let it contact its home server with a status report, which is what all IoT devices typically do. Also, keep in mind that about 95% of all domestic fires are started due to faulty electrical wiring. IoT is also meant to be about short-range Wi-Fi and radio connections. Connectivity and low price take precedence over things such as privacy and reliability, so keep in mind that IoT is not meant to be safe. IoT is actually so thoroughly insecure that you should seriously investigate how Wi-Fi and RFID work and see if you can find ways to implement custom security protocols in your household before using a single IoT device. Exercise caution when experimenting with Wi-Fi and radio waves as doing that can get you in a whole other heap of trouble.

We know from news reports that these matters are investigated by the Federal Communications Commission (FCC), who takes its job seriously. In short, FCC is like the FBI with tuning forks that hunt down everyone who might think about messing with radio waves. In October 2017, a certain Jay Peralta^[8] was fined \$400,000 for interfering with NYPD radio systems by issuing a total of nine unlawful messages over police frequencies during 2016. Jay was charged with a total of 21 counts that included terrorist threats, aggravated harassment and filing false reports, carrying a twenty-year prison sentence.

Diagnostic tools are central to becoming an ethical hacker. Being able to estimate what's happening is one skill that might serve you up to a point, but knowing what's going on gives you tremendous power because correct information is the secret to living the fullest life imaginable. Which tools to choose and how are, again, ridiculously vague questions, but you can start with baby steps and slowly build your collection based on what you discover over the course of a few years.

Laziness is a key component in the complacency of the average US consumer. IoT companies are counting on consumers being too bored or too busy to pay attention to fine print or specifics. So, stay alert and pay attention to what's happening. Always try to get your hands on raw data and interpret it yourself rather than having a laugh-track pundit do it for you. If you can keep learning and improving your ethical hacking skill, you'll be outpacing these tech companies by leaps and bounds, allowing you always to stay ten steps ahead of them.

One vigilante hacker took it upon himself to test out IoT the best way possible – by **bricking** them, which means he destroys their functionality. The hacker's nickname is Janitor, and his malware is called BrickBot, with the sole purpose of scanning the internet for insecure IoT devices with default usernames and passwords to infect and corrupt their firmware^[9], the essential code baked into the device. TVs, cameras, lightbulbs, toilets and everything else is all liable to be hit by BrickBot, and they'll all be made into expensive paperweights.

According to Janitor, 90% of all IoT cameras made by a certain manufacturer were set up with default passwords, allowing anyone to hack into them. What Janitor is doing is sadly a crime because there's still no legal consequence for having an insecure IoT device, but there is for downright destruction of property. So, what happens when IoT manufacturers start making medical implants or wearables with similar shoddy security and people start dying because of such lack of care?

What would happen if IoT devices were legally mandated to come with security warnings, such as “this device may cause a network to be hacked”? In California, every product sold already contains a warning label^[10], but the obtusely vague Proposition 65 didn't mandate companies to tell consumers what exactly is dangerous, where it's found on or in the product, why it's there for or what the actual risks are. Without any of that information, how is the warning label meant to be of any use?

With buildings having to display these same warnings too, citizens started bounty hunting and suing those companies who didn't have enough warning stickers, so companies and individuals plastered them everywhere. Nobody can tell how many warning stickers are enough to make the company shielded under the law. So, with every product and area getting a "something in here may lead to congenital disabilities/cancer", customers simply blank them out of their consciousness and do their business anyway because we have to keep living no matter what.

This oversaturation with warning to the point of indifference is a problem observed with internet users, called banner blindness and is easily explained as attention withdrawal. We all have a limited amount of attention at our disposal. When something is boring to us, it means we've determined it unworthy of our attention and want to do something else. With banner blindness, the warning or information displayed shows itself as thoroughly useless to the point people actively block out anything looking like a banner. If you're surfing the internet from an EU location, you're probably already doing this whenever you dismiss the cookie warning banner – there's no useful information in the warning banner, so you simply ignore it.

One humorous suggestion for warning labels presents "scientifically responsible"^[11] warnings, such as "this product consists of 99.9999999999% empty space" and "some quantum physics theories suggest that when the consumer is not directly observing this product, it may cease to exist or will exist only in a vague and undetermined state." While scientifically correct, these warnings would likely freak out people curious enough to read them but not curious enough to research them. This is at the core of all the warning label drama – we can't make consumers care about what they're using until it's too late.

What we can do is educate ourselves as best as we can and try to do just a little bit of good for the world. Perhaps this can take the form of tutoring individuals on the dangers of IoT, writing blog posts or just talking to those in the vicinity on the topic when the opportunity arises. The mainstream media will never consider talking about IoT unless there's some outrageous context for it – a.k.a., if it bleeds, it leads. At least kids should be encouraged to tinker with IoT and improve it whenever possible, which would teach them ethical hacking from early childhood and also give them practical, engineering skills.

Though plenty of anarchist-oriented people involved with IoT scoff at the

prospect of government regulation, the thing is that market incentives with IoT are chaotic, and there's no clear way for market forces to balance each other out within the current copyright and patent framework. Why should the producer of a cheap and widespread IoT device care if it's involved in a DDoS attack? Should the consumer care? What happens with a company's liabilities once it gets dissolved? What's to stop any given person from making a company, producing millions of IoT devices, reaping profits, shuttering the business and disappearing when customers start clamoring for security updates?

What governments can do is set soft limits on copyrights and patents in cases where the manufacturer of a device stops updating or maintaining it, perhaps even mandating that source code be made freely available in case the company is dissolved. In this way, the general public would be legally protected if it tried reverse-engineering devices and openly offering solutions to the IoT cybersecurity problem. Right now, there's no political will to think ahead; democratic elections lead to a constant rotation of elected officials, who have no incentive to offer long-term solutions, the results of which their political opponents can claim as their own. It's stupid, silly and egocentric but that's how it works and unless we upgrade our political process and our mindset, we'll be buried under a mountain of insecure IoT devices that are pretty much nobody's fault.

In one case, an IoT garage door opener manufacturer actually took revenge on an unhappy customer^[12]. On April 1, 2017, a certain Martin left a vitriolic comment on the community forum related to Garadget, cursing the iPhone app. His comment got no replies. Soon after, Garadget's Amazon page got a negative review from Martin, which prompted Garadget's developer to block his device from the cloud services it needed to operate. As is customary in such cases, other Amazon users took up pitchforks and stormed Garadget's Amazon page to show solidarity with Martin, review-bombing the product.

Analysis of the circumstances showed that the manufacturer was primarily developing apps and decided to jump into IoT to push out his app more than anything else. As such, there was no tech support for users, who had to rely on community forums or just plain experimenting with the device to make it work. The device itself relied on cloud servers to do its function, which meant it constantly phoned home. This was also a crucial vulnerability in the design as the developer could cut off any given user for any given reason from the servers. Luckily, Martin bought Garadget off of Amazon, and he

could ask for a refund but what would've happened if he'd bought it directly from the developer?

IoT devices will require a sizable expansion of the infrastructure, meaning servers to crunch the data and produce a conclusion. This will mean an expansion of existing internet address space to accommodate the burgeoning number of devices. Right now, the internet is on IPv4, which uses 32-bit numbering and provides 4,294,967,296 (2^{32}) addresses. The proposed upgrade that's meant to provide more than enough addresses is IPv6, which uses 128-bit numbering to provide 340,282,366,920,938,463,374,607,431,768,211,456 (2^{128}) addresses. The number is read as 340 *undecillions* in the US (340 *sextillions* in the rest of the world). Now there's a couple of fun words to start the day off. But, who in their right mind would need that many IP addresses?

Chapter 4 – Internet of Things

In the US, IoT spending is mostly driven by government agencies, in particular, the Department of Homeland Security and NASA, which are working on sensors on munitions and related projects. A report by Govini^[13], a company focused on providing federal agencies in the US with aggregate data needed to establish long-term policies, shows that federal spending for IoT sensors alone has nearly tripled between 2011-2015, going from \$578MM to \$1.6bn.

Military applications involve sensors on blimps to watch over supply lines and on tethered balloons to spot incoming threats. These systems were already tested in Iraq and Afghanistan, with the idea that they'll eventually be adapted for domestic and civilian use, such as to protect the US-Mexico border against illegal immigrants. NASA is working with universities to bring IoT sensors to areas such as healthcare, with products such as medical wearables that automatically test the blood sugar level of people with diabetes without drawing blood.

General Services Administration (GSA) is an independent US government body assigned to the task of managing the operation of federal agencies at a fundamental level. GSA manages some 10,000 government buildings across the nation and in 2013 decided to trial an IoT initiative called GSALink in 81 of them^[14]. 13,000 sensors were embedded in trial properties, generating 27 million data points each day. Employees cooperate with GSALink by digitally registering their workspace and working wherever suits them, with lights and air conditioning managed automatically as people move through the building. For example, if a meeting is scheduled at a certain conference room, GSALink will automatically turn on the air conditioning a few hours prior and turn it off when the meeting ends. If an employee doesn't like the air temperature, he can ask GSALink to change it, at which point the system will query nearby employees and get an average of their votes before changing anything.

While GSALink does give hope in that the federal government will finally get something done better than the civilian sector, there's looming danger from overly broad adoption of IoT before clear guidelines have been adopted. A 2017 report^[15] issued by the US Government Accountability Office lists the potential ways an IoT device could be compromised. Supply chain sabotage would come through the manufacturer of the device or its software

embedding a malicious feature that would leave the device vulnerable. At best, the device will passively collect data; at worst, it will be used to hack into networks.

Limited encryption and transmission of viewable data is another problem since a hacker wouldn't need any access to the IoT device but simply a position somewhere on the intermediary network to scoop up the data. IoT devices typically don't use encryption to save on costs and time-to-market. Poorly implemented hardware features leading to little or no cybersecurity would be another weakness of IoT. Again, IoT devices aren't meant to withstand hacking attacks or any kind of adversarial behavior.

Poor default passwords can lead to wide-scale security breaches while a lack of upgrade or patching potential could cause a situation where an IoT network isn't patchable at all. Unpatched devices will remain functional despite being outdated, just like regular computers, which will be a source of temptation just to ignore any vulnerabilities. Rogue applications can be installed on IoT devices by careless users, gathering classified or private data for the benefit of commercial companies. Once data leaves containment, there's typically no way to know who else got it and what was done with it.

IoT wearables can track the geographical location of personnel and report their location to create a detailed schematic of patrol routes or classified facilities. This already happened with Strava's smartphone app intended for runners and cyclists to measure their progress that inadvertently showed locations of facilities^[16] in Antarctica^[17] and other remote locations.

The US Department of Defense has already claimed a chunk of the entire IPv6 namespace, namely 42 decillion^[18] or about 0.01% of all IPv6 addresses. One reason for this could be security through obscurity – the idea being that all government agencies' servers can hide themselves in the proverbial haystack the size of the Solar System. Right now, it's fairly easy for any remote attacker to scan through IPv4 addresses to find entry points and just keep trying until he or she gets in, with the defender having to invest resources to secure his or her systems against intrusion actively.

Hacking attacks are precisely dangerous because it's impractical to defend against them; as new features and hardware are grafted onto the underlying infrastructure, the possibility of unintended interactions that cause a glitch or a bug increases dramatically. At some point, there comes a time for a paradigm shift, which would, in this case, be switching over to IPv6. Still,

hackers will get some spiffy tools of their own in the form of quantum computers.

Digital technology we currently use revolves around magnetism and specks of magnetic charge representing 0s and 1s on hard disks to store data. As years went by, we got better and better at storing tinier specks on hard disks, meaning we now have storage of several terabytes on a device the size of a piece of toast. However, there are some problems with this approach. The hardware we have cannot properly scale, meaning there's a soft limit as to how many computers we can stack on top of one another and expect it to keep solving bigger and bigger problems, such as being able to model Earth's atmosphere accurately.

Another issue is the hard limit as to how small a speck of magnetic charge can be before it randomly dissipates, meaning the data is suddenly lost, at which point the hard disk is no longer a reliable device. This limit represents a strong barrier against further miniaturization. Prior to asking questions we can't answer, or making devices that randomly fail in a desperate attempt to answer them, we should start thinking about a replacement for magnetism in computing, which is atomic bits of data known as "quants" or "qubits".

The quantum world is a really bizarre one. It underlies our regular world full of chairs, bottles, and chandeliers, but expected rules of cause and effect that apply to those, don't apply there. For example, in the quantum world, there's a rule known as "entanglement" that states two particles may become joined for some reason and then instantly affect one another over arbitrarily long distances. So, shaking a quantum chandelier in Tokyo may instantly break a quantum chair in Austin, Texas and fill a quantum water bottle in Antarctica. There's very little logic to how quants interact, but scientists are desperate for any sort of paradigm shift that would help them answer burning questions about the nature of the world. What could possibly go wrong?

Quantum computers are, for now, only minute experiments used to regale the public and let scientists endlessly theorize, as through IBM's Q Experience. When we do get working quantum computers, in theory, they will fit the head of a pin and be able to crack any cryptography and search through the entire IPv6 namespace in a reasonable amount of time. With such inconceivable computing power, programmers might finally be able to create an actual AI, a thinking machine brain that could fit the space of a human cranium. Its theoretical capabilities are shrouded in mystery, but it would conceivably be at least on par with its creators – if not several degrees of magnitude smarter.

Until then, we're left to use what little computing power we can cram into tiny IoT devices to produce things such as:

A smart toilet with ambient lighting and speakers for total immersion

News about Kohler's smart toilet dropped with a splash at the annual 2019 Consumer Electronics Show (CES) in Las Vegas, Nevada. Produced by a respectable plumbing company, Numi 2.0^[19] is an intelligent toilet that can set the mood lighting and respond to voice commands using Alexa for a hands-free toilet experience worthy of a king. There's a separate seat-warming add-on too, but the main gimmick is having Numi work together with smart mirrors, a smart bathtub, and a smart shower, all produced by Kohler, so you can stay squeaky clean without ever letting go of your smartphone. The price for Numi 2.0 is \$7,000 (\$9,000 for a jet-black version).

A smart fishing rod for gathering stats on the spot

Among companies angling for attention at the 2019 CES was Cyber Fishing^[20], presenting its Smart Fishing Rod to the audience. Well, what's wrong with a regular fishing rod? It can't record stats or locations of best catches, forcing the hapless fisher to write down the best spots or, God forbid, memorize them. The Smart Sensor is at the core of the Smart Fishing Rod, automatically capturing all the relevant data and letting the user easily share it online with others for fishing spots crowdsourcing. The downside is – there's no way to entertain the guys with exaggerated fishing tales.

A smart air purifier that moves around the house

Finally, you can breathe a sigh of relief through Atmobot^[21], an Ecovacs autonomous air purifier teased at CES 2019 that looks like a trash can. It moves from room to room to purify the air and vacuum the carpet, with an optional array of fixture sensors sold separately to help Atmobot detect when the air gets bad to move there on its own. Deebot is Atmobot's younger sister that simply cleans the floor with in-built support for voice controls and a detachable mopping sponge. The previous version of Deebot was priced at AU\$999 and tended to get stuck on carpets and cables^[22].

A smart water faucet with LED lighting

In the age of IoT, you can't even drink a glass of water without a smartphone app. Tern Water^[23] offers a "smart water faucet" that ties together with the app to warn the owner when the filter is about to expire or that there are contaminants in the pipes. The entire kit is \$250, but the trick is that there's a monthly subscription service for whatever reason. The tap does have a cool LED indicator but also houses a battery that lasts up to one year.

Smart scented candles with the scent of money set on fire

Moodo^[24] is a palm-sized box that fits four fragrance capsules lasting 60 hours. Of course, it's connected to a smartphone app that lets the owner adjust the aroma combination according to the app's mood meter or schedule scent release right as he or she is about to arrive home. You can also adjust Moodo halfway across the world – because why not? What's the catch? Moodo is about \$160, and each capsule costs another 35. MoodoGo is the portable version for the car that goes into the lighter port.

Smart farms

In what is most certainly bullish news, Australian farmers have been trying to incorporate IoT into their work and digitize everything from cows to fences^[25]. Unreliable connectivity and data quotas torpedoed these ambitious plans, and one of these farmers said, “We can’t apply [IoT] in 90% of situations.” A conglomerate of grazing businesses spread over 6,000 hectares, Carwool Pastoral, deployed over 200 IoT devices and found only a handful, such as smoke alarms, cattle tags, silo level monitors and shed condition monitors, had any appreciable use. Even when devices did work, the wilderness was so poorly covered with internet reception that farmers couldn’t rely on any IoT device that required the internet to function. To make matters worse each IoT device came with a separate app, and app developers had no intention of cooperating or combining their data into a single stream to make the farmers’ lives easier; imagine juggling 55 apps to check on your farmstead.

Smart sprinklers that can be turned on and off from halfway across the world

Named Sprinkl^[26], this set of IoT sprinkler modules eliminates any need for human oversight, but at what price? With Alexa-controlled interface and priced at \$225, Control SR-400 represents the sprinkler hub that requires a Wi-Fi connection to work, with the only flaw that it has to control sixteen zones or more. Sprinkl requires smart sensors Sense SR-100, one per zone and each costing \$45. Another Sprinkl module is Conserve SR-301, priced at \$79, that adds smartphone control to sprinklers, marketed as “turn off the sprinklers from your phone anywhere in the world.” The total cost for all three modules is at least \$1,024. So, is there any actual benefit to having the possibility to turn your sprinklers on and off halfway across the world?

A smart block of wood

You thought it was a joke, didn't you? It's almost zen-like with the artistic flair that defies belief – just a block of wood and the internet. Now presenting Mui, a smart home hub that has been described as “an elegant device” and is meant to mimic the woody look of furniture. It's essentially a narrow plank with electronics embedded inside that's hung on the wall and shows essential information, such as thermostat temperature. There's a display panel on Mui, but the material is actual wood coming from Japanese Hida forest. Mui already got \$115,000 out of required \$100,000 on Kickstarter^[27] and should be delivered sometime at the start of 2019, priced at \$549 for backers and \$999 in retail.

A smart padlock

Dubbed “the world’s worst padlock” and described as “it’s EVEN WORSE than we thought” by security firm Sophos^[28], Tapplock is a product of a Canadian IoT company that decided to create its own cryptographic protection scheme. Bluetooth-enabled padlock with a fingerprint scanner, Tapplock has a fatal flaw – knowing the network address of the padlock is enough to crack its cryptographic protection. Since network addresses are meant to be broadcast publicly, the result was anyone who wanted to take a look could easily crack Tapplock, which another security firm did by making a program that performed the hack in two seconds straight. Priced at \$99, Tapplock turned out to be completely hackable even remotely, without any physical access to the device, but also helpfully revealed the location of the padlock so the hacker could go waltz right in and pick up whatever valuables it was meant to guard. These vulnerabilities have been patched.

A smart menstrual cup

Small, disposable and utterly hackable, IoT devices have shown themselves insecure over and over again. So why not put one inside your own body? Looncup^[29] is literally at the bleeding edge of IoT development, as it's a smart menstrual cup that connects to a smartphone app using Bluetooth to show information, such as how full the cup is. There's a non-replaceable, non-rechargeable battery embedded in Looncup's silicone, meaning it can last a couple of months, after which it will serve as a regular menstrual cup. Additional functions include scanning the color of the blood for health problems and tracking the menstrual cycle. Like many other low-key IoT projects, Looncup began on Kickstarter.

Entrepreneurs are capitalizing on IoT hype by just making whatever seemed like a cool idea, but what about third world governments? There's an interesting aspect of IoT as it can help city officials manage affairs. So, what if we built cities with IoT in mind?

Smart cities

In a first world environment, IoT devices represent a luxury, such as smart doorbells and toilets, but in third world countries, smart-city systems made of IoT devices might actually become a fundamental part of the infrastructure. Mumbai is one of many Indian metropolises, housing 20 million residents that all want to ride one of city's 3 million cars or rickshaws to and from work. When monsoon season hits, usually lasting from June through August, Mumbai experiences a total traffic collapse.

With the help of IoT devices, such as traffic sensors, the local government can redirect traffic to less-used roads and at least try to assuage traffic jams. Other problems, such as sewage management, can also be approached with IoT to get a glimpse into major population trends. Mumbai, not San Francisco or New York, is thus en route to becoming the first smart city powered by metrics gathered through IoT. With population growth and the gradual collapse of local governments, especially when it comes to traffic, we might see more and more reliance on IoT to gather data and neural networks or AI to make relevant decisions. Simply put: humans can't handle governing themselves in such large cities.

Having IoT sensors on garbage bins can tell sanitation workers when they should go on a patrol and which route is the most efficient rather than them having a fixed schedule that might not reverberate with, say, people throwing out more garbage during holidays. In this way, waste flows into landfills and recycle plants in a steady and more controlled manner.

IoT thermostats could automatically reduce heating in those areas where nobody is spending any time, and IoT lightbulbs could automatically dim and brighten as people walk through rooms to save power. IoT traffic lights could adjust to traffic conditions ahead and IoT sensors on parking spots could interact with digital maps, such as Google Maps, to provide useful information on nearby parking spots. IoT sensors embedded in asphalt could report wear and tear before potholes form.

The major issue with this Utopian vision of a smart city that runs itself is the lack of standards among IoT manufacturers and programmers. Everyone does things his or her way, and it's as if apps are intentionally made to conflict with one another. Lack of standards is a common sore point when it comes to software, with each developer stubbornly refusing to adapt or cooperate with his or her competition. As software companies rise to prominence and wane

away, the consumer market is peppered with proprietary formats and designs that become useless after official support ends.

It took us decades of anguished struggle between software developers to arrive at having PDF and MP3 formats that are universally recognized across all mobile and desktop devices; intermediary file formats were used for a little bit and quickly forgotten. Applied to IoT, this means we can likely expect an explosion in experimentation with regards to designs and formats until the entire world settles on a couple of solid standards after a few decades. Of course, early adopters will get shafted, but the prospects of a smart city are so tantalizing.

Ideally, one overarching AI could finely control the environment throughout the entire city, cutting down on passive losses common to all power and heating distribution systems. A fleet of drones could be sent out automatically to clean up and scrub after a festival or public party was held. By analyzing behavior patterns, AI could know who's about to get sick and what's the effective treatment *before* the person feels the first symptoms. IoT devices in our homes would act as alarms and notifiers – ever had that overwhelming fear you left the stove on? IoT could be connected to heat sensors in appliances that would buzz your smartphone if you actually did.

That's the ideal future, but what we're looking at is mostly a chaotic array of gadgets cool in design but barely functional because they refuse to work together. If Samsung does waste disposal, Google does heating, and Apple does power supply, the three might intentionally sabotage one another, but if Samsung does *all three* in one city, now we're getting the coherence IoT gadgets need to function as a part of a broader network.

What are the implications for the democratic process in one such city where a company essentially has insight in all consumption and behavior patterns of the population? Can you criticize Samsung in Samsungville? How trustworthy are elected officials that kowtow to tech giants? Can we even escape the influence of tech giants?

Chapter 5 – Under The Cushy Foot of Tech Giants

Those who tried cutting tech giants out of their personal and work life experienced a swift and humiliating defeat. In January 2019, a Gizmodo journalist tried to cut out five tech giants for a week each, with the third week dedicated to Google^[30]. After setting up VPN filters that blocked Google's IPs for her devices, the journalist was essentially unable to do any work and couldn't even use Uber or Lyft, as they engage Google Maps to operate correctly. Spotify hosts its content on Google's cloud, so no music either; AirBnB dumps its photos there as well, so they didn't load. Fonts, analytics and just random tidbits of code are all helpfully hosted by Google, which means the websites she used were acting up in all sorts of unusual ways. In short, her browsing experience was like being back in the early 90s: slow, messy and barely usable.

Eventually, we might all be living under a cushy foot of tech giants to the point that trying to get out will be too uncomfortable. Bit by bit, who we are and what we do is being pinned down to the point where we can't keep any secrets, shifting our life to the cloud where tech companies have supreme control of our activities. Arcane rules of conduct and vague terms of service already dominate on social media platforms, where accounts can get suspended for satirical content that others share freely. In some cases, tech companies can even decide to enforce rules of conduct *off their platform*.

In late 2018, the crowdfunding platform Patreon decided to ban a user^[31], a certain Carl Benjamin, for comments he made on YouTube, despite Patreon's terms of service never mentioning or implying that option. This would apparently open Patreon to lawsuits for breach of contract, which is what terms of service legally represent, but so far none of those have materialized. Instead, users voted with their feet and left the website in search of better options only to find that, well, there aren't any. So, Patreon is a minnow website dedicated to a niche audience. What does one do if Google decides to enforce its terms of service for an off-hand comment made in real life and comprehensively de-platform an individual? It's actually impossible to punish Google in any significant way without destroying the internet as we know it.

There are simply no enshrined rights for individuals online like there are in the real world through constitutional charters, yet we sorely need to have them formulated. The same way our societies went through tumultuous times

where whimsical tyrants lashed out at citizenry before evolving to democratic cultures that respect human dignity, we need to have our online world evolve *before* IoT becomes a staple. We needed to have them *yesterday*, or technological progress will become just another outlet for whimsical tyrants to feed their ego. Power truly does corrupt, and when tech giants wield so much power that they can undo a person from the internet, they kind of lose sight of reality.

Google, Amazon and other tech giants thus become 500-pound gorillas that sit in our home and occasionally lend a hand to do the heavy lifting – just keep your head down and don't look them in the eye or they'll kick you out. When other, smaller gorillas see we're not putting up a fight, they'll move in too, and our home will become a circus rather than a sanctified resting spot. There doesn't seem to be any pushback to this encroachment, and the media is disinterested in headlines that involve a thoughtful debate about such an abstract issue; today's media is all about clickbait, meaning titles that attract clicks and push copies at newsstands.

The best protection we can employ is to fragment our online persona as much as possible, meaning we should separate all our accounts so banning any of them doesn't collapse our entire digital identity. So, you could be "Karl" on YouTube, Karl123 on Yahoo, Karl234 on Amazon and so on. The more you can separate between your online accounts, the more you are shielded from the tyranny of tech companies. Yes, this does mean remembering a dozen or so very strong passwords and constantly logging in and out of accounts, but that's the price of freedom of speech and thought. If possible, try to have a dedicated device for each service: a smartphone only for YouTube, a laptop only for Yahoo, a tablet only for Amazon, and so on.

The trite argument of "just don't use the website if you don't like it" sort of works because the individual can still go in the real world and do something else. What happens when IoT makes everything connected to the internet, and there's no escape from denial of service? Fine, so we'll make an alternative to Google – we'll make a company that cherishes freedom of speech and honors terms of service. Any competitor to Google can be blocked from Google services, which would include marking that person's website as malicious, removing his or her app from Google Play, removing related search results from Google Search and denying him or her access to Google Cloud storage. It's effectively as if being un-personed and completely erased from existence in the digital age.

When tech giants go mad gorilla against one another is when we pull out some popcorn and enjoy the carnage. January 29, 2019, is when Apple discovered that Facebook was abusing its enterprise privileges and blocked all Facebook apps from functioning off of Apple's cloud for two days^[32]. Facebook employees experienced the same effect as that Gizmodo journalist – they couldn't even check their calendars or schedule a lunch because all of their infrastructure was hosted on Apple cloud.

What happened was that Facebook enlisted teenagers, as young as thirteen, through ads and had them install a **VPN**, a virtual private network that filters traffic, using an app that intercepted and scanned their internet traffic for some \$20 in gift cards. This was against all imaginable privacy policies and was likely illegal too. In June 2018, one Facebook app, Onavo Protect, that did the same thing and was already banned by Apple, then updated its developer policy to stop such abuse from happening again. This time, the app was codenamed Project Atlas, but journalists that peeked inside the app's code found numerous references to Onavo Protect, confirming that the spying project was simply renamed.

For two days, the Apple cloud was inaccessible to Facebook, making much of its internal work impossible. After that, Facebook was allowed back in, despite doing what would have made some other developer permanently banned from Apple cloud. That's simply the kind of influence tech giants wield that allows them to create their bubble of sovereignty where they can ignore laws and bypass all rules and morals. So, what happens if an upstart company decides to offer an alternative? It simply gets bought off.

Chapter 6 – The Power of Infinite Funds

Onavo Protect was initially a VPN app that helped users stop ads and trackers that hogged their bandwidth, but Facebook acquired the company in 2014 for some \$120MM. Because user data is considered an asset of the company, by acquiring Onavo, Facebook got all the data that users thought would be kept private. Seeing how Facebook made some \$50bn in 2018, \$120MM was chump change, but Onavo executives were most likely ecstatic because they got to live the dream of cashing out and spending their retirement in the Bahamas. There was no way for Onavo to resist the buy-off since executives are supposed to do whatever isn't strictly illegal that will make the company money; had a moral Onavo executive decided he didn't want to sell, the rest of the board would have had to relieve him prior to selling the company anyway.

Buying Onavo let Facebook essentially spy on its users, which showed a certain app was being used much more than Facebook's Messenger – WhatsApp. Facebook jumped in and bought WhatsApp for \$19bn, which was and still is the largest acquisition of a company in history. The motivation for buying WhatsApp was to create a sort of umbrella network of free essential services, such as Facebook for organizing social events, WhatsApp for messaging, and so on. Dubbed "internet.org", the app intended to deliver these services was named Free Basics in 2015 and was aimed at the poorest nations, India in particular. It's too bad India banned Free Basics, stating it had discriminatory tariffs.

So, if Facebook offered free services, how was it going to foot the bill? With users' private data. By making itself the gatekeeper of internet services for people too poor to afford a non-spying alternative, Facebook wanted to reach an unprecedented level of insight and control over its users, who would have nowhere else to go. The trick is to rope in children as early as possible and hook them into using social services through Facebook's funnel, which is why Project Atlas was intentionally aimed at the youngest teens possible – a company that can hook kids into its ecosystem of products basically has them for life. This is the strategy of soda makers, breakfast cereal producers and other such companies that deal with products that provide immediate gratification.

There are no protections that would aim specifically at protecting private data of children, who often have no concept of social norms and tend to say stupid

things. Despite everyone taking up arms whenever children are endangered, this doesn't seem to apply to persistent, low-key threats but only dramatic and explosive ones, as explained in "Freakonomics"^[33]; children are constantly warned about the dangers of guns but never about the dangers of swimming pools, which are statistically much deadlier but aren't as violent.

Prior to the internet becoming widespread, children had the luxury to be obnoxious and inane, with only their immediate surroundings knowing about it. Thanks to social media, their stupidity can now follow them their entire life, impacting their employment and relationships decades in the future. Of course, companies have not an inkling of concern about this – since they only want to make money and "make the world a better place". In any case, the initial discovery that Facebook was using Onavo Protect to leech data was what prompted Apple to tighten up its developer privacy policy. Google was found using a similar scheme^[34] to leech data from teens but abruptly ended the program when journalists asked about it. When the same corporate logic is applied to toys is when things take a more sinister turn.

Chapter 7 – IoT Toys

When applied to toys, IoT becomes a terrifying tool of surveillance, even when done unintentionally and out of negligence. Children's most intimate moments with their parents become open to any scoundrel who might want to listen in to their conversations or scammers who might want to blackmail the company. In one such case, IoT teddy bears were found to be storing all recorded messages and profile info in a public-facing database.

Marketed as “A Message You Can Hug,” CloudPets is a collection of teddy bears and other plush toys that can record and transmit voice messages between kids and their parents. The only problem was – all personal data related to CloudPets was stored in a public database that was accessible online by anyone who wanted to spend a few minutes looking for it. In February 2017, word got out to Troy Hunt, an Australian cybersecurity researcher who at one point testified in front of the US Congress on data breaches, so he took a peek^[35] and was stunned at how bad the security was.

CloudPets had over 2 million voice messages belonging to about 820,000 owners stored online to cut on costs and engineering complexity needed to make the toy itself store messages, which would be the safest thing to do. E-mail addresses and encrypted passwords were stored in MongoDB, an open source database that uses an easily copyable text format rather than rows and columns like Excel. Contacting the owner of CloudPets via e-mail about the vulnerability produced no reply; apparently, there was nobody at the wheel.

Troy then analyzed CloudPets' app behavior and discovered that it stored profiles on Amazon servers, containing profile photos, names of kids, dates of birth and their relationship to the adult that bought them the toy. Recordings of kids' messages could be accessed by simply knowing the file path to them on the server, which Troy tested out and actually heard a few of them. Weak passwords were another headache, with the official CloudPets tutorial showing account creation that included using a mere three-character password. Troy tested the stored passwords using a dictionary attack and found plenty of them being “12345”, “password” and other such easily hackable ones.

In other words, there was no requirement for a password to be of any particular length or complexity; one could put a single character as a CloudPets password. Once a password is found, there's an associated e-mail address displayed right next to it in the database, so the hacker can simply log

in as a legitimate user to the given CloudPets account.

Hackers didn't stop merely at listening to lovely audio messages; they actually deleted the exposed databases and replaced them with ransom messages stating that a backup could be retrieved for one Bitcoin, which was about \$1,000 at the time. An e-mail address attached to the ransom note indicated the hackers were from India, but none of that mattered because there was still no official reply. The owner company experienced a 99% drop in stock price since releasing CloudPets, resulting in downsizing that likely took out all tech support and interest in maintaining the product line.

The reply by Californian-based CloudPets CEO, Mark Meyers, was that those voice recordings were not stolen, that headlines of 2 million messages being leaked online were false, and that the security weaknesses were a "very minimal issue." Technically, he is correct, but what about addressing the underlying problems? When asked about not having password strength requirements, Mark replied, "How much is too much?" When asked about the warnings, Mark said that it's their policy to ignore warnings coming from random people, indicating that a security issue is only a big deal if it's about to appear in nationwide publications.

Troy's advice with buying and using IoT toys is to "assume breach" – meaning that you should consider security non-existent unless proven otherwise. This means that it's up to you to discover and change default settings (if any), set a strong password, and otherwise discover how the device works before handing it off to the kid or merely bringing it inside your home. For busy parents who think IoT toys will help them save time, it's quite the contrary – now they've got to learn the essentials of cybersecurity and networking or risk all their personal information being exposed to unsavory characters.

CloudPets' story is one of few where someone with expertise cared enough to take a look, analyze the situation, and write up a report; otherwise, none of it would have ever come out to the public. For CEOs that produce IoT toys, each sale is relevant insofar as it generates revenue; products and services become a liability that needs to be minimized and ignored whenever possible. CloudPets' CEO response is the perfect example of how to slyly handle a security incident with rhetorical questions, dismissals, and vague explanations. Unless we as consumers start taking care with ours and our children's security in a world filled with IoT, nobody will.

In the hands of hackers, all IoT devices become toys, as it were. The more we scale up IoT without addressing these fundamental security issues, the more we'll be exposing ourselves and our loved ones to endangerment. It's quite possible that the way we imagine IoT right now isn't feasible on a worldwide scale, but can only work locally. We're simply stumped for answers. Perhaps the solution is to wear personalized IoT sensors and have IoT access points inside a house or workplace reading movement, speech, mood, and behavior. But again – can it scale?

The same way a tiny flea can jump a foot in the air but wouldn't be able to move if it were scaled up because the materials it's made out of wouldn't be strong enough, some systems only work when done on a tiny scale, and it appears IoT is one of these. Fragmentation and localization of IoT networks take the wind out of the sails of most IoT marketing strategies that would ideally want everything to become IoT – of course, with proprietary hardware and software. So, if IoT vendors use third-party open standard software, what's the risk a single exploit or bug can bring the entire network down?

Security firm Senrio conceptualized an attack called **Devil's Ivy**, where a bug found in the third-party code used in IoT devices can be exploited to compromise devices themselves. With Devil's Ivy, the focus is on gSOAP, a C++ framework used to develop **SOAP**, Simple Object Access Protocol, which is meant to unify different IoT operating systems and frameworks through XML. Basically, two IoT devices using SOAP can communicate as long as there's an internet connection between them no matter what their manufacturer or operating system. It sounds great – for hackers that is.

Devil's Ivy allows a hacker to exploit a bug in SOAP found on an IoT device, such as a CCTV camera, by sending a malicious 2GB payload that overflows the device and resets it to factory settings. Then, the hacker logs into the device and can either try to access other vulnerable devices or just quietly absorb information by looking over the shoulders of employees as they're typing in their passwords. If a router is hacked this way, the hacker just got access to all network traffic, which is as good as full access to all devices on that network. By the time the owners spot a problem, *and* the device vendor issues a patch, hackers could've been spying on people for decades.

The company that made SOAP, Genivia, issued a patch for that particular bug while noting at least 34 IoT vendors use SOAP. What happens if a commonly used IoT framework is found weak at some point in time, but there's no central authority to issue a patch? Who can force companies to

patch their IoT products if that comes at a cost? Right now, the only advocates for IoT standards are non-profits.

The Open Connectivity Foundation^[36] (OCF) is a valiant effort to create IoT security and connectivity standards by advocating for “security by design”, meaning hardware and software is built in expectation of being hacked. Pushing for this concept is in a sense a way to preemptively quash all complaints from companies that security is expensive; by baking security into initial design, all costs are offset onto customers. OCF also aims to create a model for centralized IoT management using public keys infrastructure, meaning a company can update all its IoT products in one go.

It sounds great, and it makes a juicy article headline, but it’s unlikely to happen. Companies are typically prone to what’s known as “virtue signaling”, where executives say whatever makes them look good but do the most selfish thing anyway. This is why every Silicon Valley executive, such as Apple’s Tim Cook, spouts exuberant phrases in the vein of “Apple makes the world a better place”^[37]; Apple will then allow Facebook to collect teenager data using Apple Cloud. If tech giants are dominating the digital world, at least we have the real world – where we can go outside and enjoy nature without IoT, right?

Chapter 8 – Bio-robotics

It's one thing to be surrounded by IoT indoors because you can simply step out for some privacy, but it's something quite different when IoT gets deployed outdoors as well. If there's a commercial application to outdoors IoT deployment, you can be sure that it's going to get done one way or another.

IoT devices enable real-time collection of granular data to visualize complex environments through the mass movement of individual units. One such IoT concept imagines bumblebees as “living platforms” for IoT sensors. In “Living IoT: A Flying Wireless Platform on Live Insects”^[38], five researchers from the University of Washington imagined a world where bumblebees equipped with IoT sensors help humans do what they termed **bio-robotics**, grafting of digital devices on analog beings. It's actually quite a practical approach that solves many insurmountable problems.

Engineering a drone to serve as a mobile IoT platform comes with many headaches, such as design, energy source, wingspan, miniaturization and getting funds for all that. Living insects are an example of living perfection, so simply tacking lightweight IoT sensors on them essentially hacks all those problems in one fell swoop. Bumblebees are extremely efficient in gathering resources and spending as few of them as possible for their operations, again making them a great surrogate for IoT drones. They also have navigation down to a T with innate magnetic sensors that help them align with the Earth's magnetic field, hence why we say “beeline” for a perfectly straight line.

Weighing 0.1g, the sensor carried by the bumblebee can be detected within 80m of the nearest access point and can transmit 1kbps of data when the insect is back at the hive. The battery on board the bumblebee lasts up to seven hours while recording the current location fifteen times a minute. With alternative sensor modules for humidity, temperature, and light, a swarm of bumblebees can be turned into an efficient mapping tool to find out the most suitable location for any given plant that requires a certain combination of humidity, temperature, and light to thrive.

A suggestion presented in the paper is surgically inserting IoT sensors into bumblebees at various stages of their development. The electronic waste would become an issue, but the IoT sensors could be built to transmit the location as the insect is expiring and biodegradable sensors are a possibility.

Conceivably, there could be a fleet of slightly larger IoT drones designed to go around and clean up dead IoT bumblebees, a fleet of even larger drones for picking those up, and so on. It would be drones all the way up, but would that make the world a better place?

The main limitation in IoT drone manufacturing is battery life. Currently, the most efficient batteries use lithium, which is thought to be the lightest material possible for storing energy. As scientists tweak the lithium formula, it's possible they'll find ways to squeeze a bit more power out of battery design, but then the progress inevitably plateaus; unless there's a paradigm shift past that, we'll be stuck with lithium batteries for the rest of the future. In comparison, bumblebees use macronutrients – proteins, fats, and carbohydrates – to power their operations. We use these three as well, simply because they're everywhere and provide lots of energy, just like fossil fuels.

As tempting as it may sound to ditch organic energy sources, we'll be stuck with them for the foreseeable future because they're relatively cheap, reliable and efficient, but they miniaturize poorly. Even electric cars run into the same problem of having to cart around bulky lithium batteries that need to recharge, which is mostly done by using the power created by burning fossil fuels. Right now, all IoT devices are merely stationary and semi-stationary appliances, so can we make IoT *wearable*? Apparently, the answer to that is a resounding “yes,” and we might actually need wearable IoT to maintain our health and social well-being.

For children with autism, daily life is a constant struggle as they can regularly fail to do the most basic things, such as tying their shoelaces. Autism itself is a catch-all term for brain dysfunction that results in socially unacceptable behavior; there's no real boundary as to what is autistic, so the illness is seen on a spectrum. Such children lack a healthy emotional response to natural circumstances, such as feeling fear when they're in danger or speaking up when they're hurt and often find themselves ostracized. Because, normally, emotions would guide us to a better life, autistic children can experience tremendous stress that leads them to bottle it all up and then have emotional meltdowns for apparently no reason. In any case, IoT wearables can really help children with autism get a grip on reality, but especially help their parent keep their sanity.

Now here's where Google Glass and the cohort of analytic eyewear could help – this is worn by people who truly can't discern emotions on the faces of people they're interacting with, such as autistic kids talking to their parents.

Because the environment is trusted and there's actual productivity lost due to inability to see emotions, there's a strong incentive to use IoT. However, what do we do if the parent can't understand their autistic child either?

In this case, an IoT bracelet serves to track the child's reaction to the environment and report the findings, perhaps to a parent's smartphone app. By monitoring the heart rate, a parent can see when the child is experiencing stress and can react straight away instead of waiting for the meltdown. How will the parent know what to do? Whatever causes the child's heart rate to slow down is what it takes to calm the kid. These bracelets are similar to those worn by fitness enthusiasts: simple, sturdy and effective.

Smart bracelets can be used by people with diabetes to scan for blood sugar levels and determine if there's any need for a reaction without drawing blood. Diabetes is so insidious that it takes an entire network of people to keep just one person with diabetes in check. Nutritionists and doctors could have instant access to this information to schedule a checkup, especially in cases where the patient has trouble moving – which is again another common symptom of diabetes. Family members could also keep an eye out for blood sugar levels and intervene when needed with a meal or just a glass of orange juice. One such bracelet is already on the market, dubbed GlucoSentry^[39]. Apple is already capitalizing on the medical wearables market with its smartwatch that can track health signals and possibly report them to a medical professional.

As technology advances, we on average have a longer lifespan, which means more chronic health problems, such as cancer and diabetes. Medicine is suffering tremendous costs with simply keeping these patients alive, so IoT wearables could help cut down costs without just hiring more people to do costly medical exams. In medicine, IoT could be indispensable because we're already struggling with the lack of specialized staff to draw blood, measure blood pressure, and so on. A doctor needs to have immediate results to make the right call or risk being sued for malpractice, but there's the lack of staff, patient discomfort, and paperwork – IoT devices neatly fix all those problems and might even administer drugs automatically. For example, MiniMed 530G is an implantable pancreas that comes with an external sensor that shows blood sugar and insulin levels.

At the 2018 Medical Sensors Design conference held in San Jose, California, projected growth of medical wearables was said to be about \$12bn by 2021. This includes masks^[40] to exercise the facial muscles that might have

degraded due to surgery or other causes but also monitors inserted in a woman's body to track ovulation.

OvulaRing is a smart ovulation monitor that tracks the core body temperature to report when an egg is released. It's already available in the European Union at \$520 for a twelve-month package.

HealthPatch MD is a wearable sensor that looks like a nicotine patch. It detects body posture, heart rate, skin temperature, and respiratory rate while triggering an alert if the wearer too suddenly changes posture, which could indicate they've fallen.

Zio XT Patch is a heart rhythm monitoring patch that can be worn up to two weeks continuously, revealing abnormal heart activity patterns. Data is sent over to the app that then crunches it using algorithms.

Quell is a knee band that looks to reduce pain and discomfort, coordinating with the smartphone app via Bluetooth to squeeze and release as needed to massage the area.

WristOx2 is a watch-like wristband that monitors oxygen saturation in the blood and changes in blood volume in the skin.

The biggest obstacle to IoT medical wearables in the US is getting FDA approval, which can take years and hundreds of thousands of dollars in clinical trials. This is why Apple released a "smartwatch" that was met with ridicule for being too simplistic for its cost; Apple exploited a loophole that allowed it to sell medical tracking technology without FDA authorization under the guise of selling a watch. It's actually a brilliant business move. The big companies know how to monetize their fanbases, but they'll also be getting a wealth of data to analyze consumer behavior on a grand scale and predict their decisions, which is termed **predictive analytics**.

Chapter 9 – Predictive Analytics

Now that a company has access to data from sources such as IoT wearables, how can it be used for analysis? First, the data is *anonymized*, which means any personally identifiable information is removed. This neatly dodges wiretapping regulation but otherwise doesn't mean much because each data source is labeled with a number that can always be traced back to the person. By gathering all the data from as many different sources as possible, the algorithm or neural network at the other side weaves a digital doppelganger to the person and tries to predict their behavior. Wait, are we so predictable?

We all have an extremely basic part of the brain called the **limbic system** that houses essential urges and functions, such as hunger, aggression, and territoriality. It is extremely *fast* and efficient, acting before we get the chance to think about it consciously. When we observe the action of the limbic system, we often do it in hindsight and can't cope with the fact it's basically working on autopilot, so we assign it ulterior motives which aren't correct in the slightest, "I meant to do that all along!" However, a machine is a dispassionate observer and can see through the fog of rationalization to deliver the *truth* that can be seen by the entire world.

Companies generally keep the findings of their research locked away, but for the first time, IoT provides everyone with a chance to participate in this grand experiment *and* get the results. While it does feel like machine spies are violating our privacy, predictive analytics can help us understand our brain because *we have no idea how the human brain works*. In this way, machines reveal our behavior to us, and companies are already capitalizing on this concept. Simply put: we're the smartest creatures on the planet and can adapt to anything as long as we can see our behavior in an objective way, such as by using IoT wearable sensors.

Humanyze^[41] touts itself as "science-backed analytics to improve your compliance." The team adjacencies metric gives each employee a rating on how good he or she is when communicating with his or her immediate team and the rest of the company. Both volume and gaps are taken into account to assess the communication risk each employee poses when it comes to the delivery of valuable work and expressed as percentages and hours.

Time allocation metric shows how each team member spends his or her time using a communication medium, broken down as chat, e-mail, phone call, and meeting. After work hours are included too, showing who's available

after they leave work and revealing cultural differences that can be grafted onto other teams if needed. Periods go from one week to one year, but ‘role’ is the most important factor in assessing appropriateness. In the end, this metric can answer that one burning question, “Are we having too many meetings?”

Communication by gender metric measures the volume and type of communication sent by each employee and broken down by gender. In this way, gendered preferences to using any particular mode of communication are revealed, helping with better integration in the company – for example, if men or women are invited more or less often to meetings or called over the phone. The result of these metrics is a 2D graph that looks like a spider web and visualizes connections between team members.

Founded in 2010 by MIT students and one professor, Humanyze utilizes sociometric ID badges to track movement and performance, essentially giving employees wearable sensors that don’t intrude on the communication or impart their opinion of the interactions taking place. Their 2008 paper “Understanding Organizational Behavior with Wearable Sensing Technology”^[42] goes into greater detail on how each personal relationship consists of four basic behaviors that fundamentally predict productivity.

The badge form was chosen since employees are often already asked to wear ID badges; Humanyze badges were pimped out with microphones, infrared transceivers and accelerometers to show movement and speech patterns. Sturdy, easy to use and unobtrusive, the badge could recognize:

- if the wearer was sitting, walking, standing or running in real time
- analyze vocal tonality shifts to measure excitement and interjections while ignoring the words themselves (again, note the sly dodging of wiretapping laws)
- the position of each wearer through triangulation of the badge’s position, with error as low as five feet (1.5 meters)
- nearby Bluetooth-enabled devices and communicate with them
- face-to-face interactions because the infrared sensors in badges could spot each other

One of the goals of the paper was to determine *interdependence*, meaning how much employees have to communicate with one another to complete any given task; if they are interdependent more than their communication skills

allow, their productivity will suffer. So, the students and their professor went to a German bank with their badges and started gathering data.

Over the course of one month, 22 employees in the bank distributed as four teams, two mid-level managers and one high-level manager wore Humanyze badges while they were on the clock, resulting in 2,200 hours of data. Individual and group performance satisfaction was also measured through a survey at the end of each workday and e-mail logs were gathered as well. Employees were 50-50 men and women, though all managers were men. Employees were split across two floors, which was another reason why the bank was interested in Humanyze's view of things – does that layout impact our performance?

The results showed that the amount of time spent with other people negatively correlated with e-mail activity, which showed that e-mail *is not* a replacement for face-to-face contact. Further, the total amount of communication *lowered* the employee's satisfaction value, as self-reported through the daily survey. Bank layout did not impact employees negatively since only the managers interacted between floors and the key premise was that an employee that is central to an organization experiences lower satisfaction. Next up, the team visited a Chicago data server configuration firm.

Conditions of the second experiment were similar: one month and 23 employees wearing badges until 1,900 hours of data were collected. All were men but their skill levels varied. Their job consisted of waiting until a field salesman contacted them with client preferences regarding a computer configuration, at which point they would use a certain program to create that configuration and send it alongside the price estimate back to the salesman.

Four types of behavior were spotted: low/high physical activity with/without speaking. High physical activity was essentially fidgeting, which is the activation of the limbic system that wants to fight or flee when it feels cornered. Frequent speaking is another sign of the same limbic system activation. Both fidgeting and frequent speaking are indicators of stress, which is known to impact productivity negatively. So, researchers theorized that those employees with the least physical activity *and* who spent the least time speaking would be the most productive.

Findings confirmed that theory and showed that the low physical activity without speaking group did their tasks 63% *faster* than the high physical

activity with speaking group. The number of follow-ups, which are repeated calls from the same salesman regarding the same configuration, was 28% *lower* in the former group compared to the latter, implying the tasks were also done more precisely. The conclusion was that “environmental distraction in an individual may trigger bursts of activity, and this distraction subsequently lowers performance.” The lesson from this would be – if you want to be efficient at your job, stop fidgeting and stay silent.

In 2011, Humanyze produced glasses that serve as a social cue detector^[43]. When worn and looking at a person’s face, these glasses use an attached headphone to explain what the other person is feeling, “Bored, disappointed”. There is even a tiny traffic light embedded on the glasses frame that warns when the other person is about to speak, so the wearer doesn’t interrupt. 24 facial points are analyzed to reach a conclusion on what the speaker is feeling.

These glasses were originally meant to help autistic people, who often have trouble reading social cues from faces, but the team that made them was stunned to find out non-autists could interpret only about 54% of emotions on faces, which is slightly better than just flipping a coin. How well do the glasses do? 64%. Companies that produce adverts or movies are already clamoring for these glasses to find out the impact their content is having because they realized *people don’t know what they’re feeling*.

To recognize our own emotions and those of others, we can pay attention to what is known as “honest signals”. These include things such as gesture mirroring, where we involuntarily repeat what the other person is doing, for example rubbing the forehead. We still respond to gestures from others; it’s just that we do it involuntarily – the gist of this is that we must become aware of our behaviors before corporations do or they’ll find a way to harness our limbic systems to their favor, such as by finding out what we truly feel.

Converus EyeDetect^[44] is another gadget that promises to replace the lie detector contraption we see in thrillers by checking the involuntary reaction of the eyes. Lie detectors work on a similar principle – we instinctively want to tell the truth but blocking that instinct shows up as increased heart rate, blood pressure, and sweating. By the way, lie detectors are not considered foolproof and are, at best, slightly better than a coin flip at producing evidence. The motive for using a lie detector is that an evildoer will reject it, thus implying he or she has something to hide – unless he or she is a psychopath that doesn’t care about the truth at all. Converus EyeDetect also

presumes that pupils react to lying the same way the rest of the body does.

So, these wearable sensors produce a tremendous quantity of data; Converus EyeDetect captures 60 data points per second per eye. How are companies meant to sift through them and find meaning? It's through creating digital brains that are trained on very simple tasks and perhaps have the intelligence of a snail but work a million times faster. After solving billions and billions of the same tasks within a day, the digital snail's brain is almost perfect in instantly finding solutions to similar problems it was trained on. This is what's known as **machine learning**.

Chapter 10 – Machine Learning

If a plant seed is placed in the dark and there's even a hint of sunlight, the plant will grow, twist and contort itself as much as needed to reach the light. If a seed is placed in a dark maze and the plant needs to solve the maze to reach the light, it will do that as well^[45]. We can arbitrarily scale the maze up, and the plant will keep struggling to find the exit, sending offshoots down separate paths for information on how to reach the light. How about placing a potted plant next to a window and rotating it away from the light to see what happens? The plant will *slowly* rotate itself back around, so the leaves absorb the most sunlight^[46]. This happens imperceptibly and, with the exception of sunflowers, we scantily notice that plants can turn around *and that they have a preference for the direction they're facing*.

How about a slime mold placed in a maze with a piece of food at the center? If the slime solves the maze, it gets a tasty treat, which it invariably does^[47], and again we can scale the maze up, and it will always get solved in the most energy-efficient manner. Fungi, mice, birds, cats, dogs, elephants and chimpanzees – every living creature shows the same innate propensity for solving spatial challenges to reach food, and even humans plopped down in the middle of a shopping mall with a grocery list will eventually amble their way out the door with a loaded cart. All creatures except **robots**, thinking machine servants.

The pride and joy of human creation, the pinnacle of mechanical engineering, and yet, robots are as dumb as rocks and can't do *anything* unless specifically instructed to do so through **code**, a set of machine-readable instructions. Any change in the environment invalidates previously written computer code; any conflict in the code leads to unpredictable behavior, which is what we call "bugs". While living creatures have the genetic code to guide them through life's challenges and mazes, robots and computers have nothing of the sort unless someone writes out a specific set of commands: if A, do B unless C. This means a robot has to have a specific code written out for every given maze, and the code needs updating whenever the maze changes or the robot is moved slightly from its starting position, or there's any change in the environment whatsoever.

Machine learning is the brilliant idea that, since living creatures have the genetic code that holds instructions and computers have code too, perhaps creating such machines that can randomly mutate their programming can lead

to something intelligent, the same way millions of years of evolution led to slimes and plants solving mazes in search for food. So far, there's been just enough progress to whet the appetites of scientists working on the concept, but there's no way to break through the conceptual barrier and create actual, independent intelligence. It's tantalizingly close and yet it appears reaching actual artificial intelligence could be the undoing of us all.

Machine learning can be used to *mimic* the actions of living intelligence to an extent, being primarily used to create **neural networks**, decentralized consensus data processors. All right, let's take a moment to unravel that conglomerate of buzzwords. The entire field of machine learning is like that, filled to the brim with hopelessly complicated expressions. In this case, decentralized means damage or corruption of any given node won't collapse the network, giving it resilience just like the one living beings have. Better yet, the network can learn to recognize and route around damaged or corrupted parts, building new structures on top of these unused subsystems. Remind you of anything? That's the *way scar tissue forms after damage to living organs*.

Neural networks would employ a consensus protocol, which means a piece of data would be flowing through the nodes in one direction and they'd all get to vote on it. If there are conflicting or obviously wrong votes contrary to what the network creator set down as ground truths, the network as a whole can reach a consensus to ignore those votes or give them lesser weight as time goes on, just like humans do. We give greater weight to information coming from certain sources we trust, though we tend to go to the other extreme and trust the minority of sources too much at the expense of listening to what the majority is saying to stay in touch with reality.

When you think about it, there's a reason why democracy is used throughout most of the world – when everyone gets to cast a vote, no matter what it is, the total of it more often than not reflects reality, which is termed **wisdom of the crowd**^[48]. We reserve the voting process for electing government officials, but neural networks allow us to crowdsource answers on any given topic. Methods of governance evolved over the course of millennia, but neural networks can find not just the best answer but the *best method for finding the best answer* in a matter of days, though they still need some help setting up the data sets.

Data would enter neural networks from trusted sets, which can be labeled or unlabeled. A small subset of data is usually used for training the neural

network and allowing it to develop fully. Thanks to cyberspace existing in a non-physical environment, data can have as much as 200 dimensions, allowing the neural network to contextualize abstract notions such as words and humor. Anything that can be perceived by a human can also be understood by a neural network and at a much faster pace. These things are perfect tools for crunching data.

Finally, neural networks are processors, meaning they produce something fundamentally novel, a result that was unknown to the network creator before the experiment started. This can range from spotting unique patterns related to cancer treatment in millions of patient data points to optimizing existing solutions related to things such as power grid management. In a very narrow sense, the neural network *creates*, which means it could technically claim copyright if it gained consciousness.

Another thing about smart machines is that there are absolutely no safeguards for them *or us* when they eventually emerge as citizens in their own right. Every societal issue we're currently struggling with, such as gender and ethnicity, will become a hundred times more complicated when smart machines enter the fray and they'll just keep evolving until reaching human levels of intelligence aka, **artificial intelligence**.

Chapter 11 – Artificial Intelligence

Intelligence can be loosely defined as the “ability to adapt to the environment” and is a remarkable predictor of survivability – a predator that can outsmart its prey can grow bigger and have more offspring and vice versa. Intelligence is always balanced out by the necessity to deal with the real world in the here and now, meaning that a cat is as smart as it needs to be to use its body effectively; any smarter than that makes it exhibit bizarre, un-catlike behavior. So, animals in nature experience a tight bond between their intelligence and the ability to deal with the real world. The two slowly evolve over millions of years, inching forward in lockstep.

Humans are the most advanced species on the planet because they can change the environment to suit their needs, perfecting both their intelligence and physical ability. We also produce tools and technology to become more comfortable and productive, making sure always to balance the two. For example, an AC unit cools the room down when it’s hot outside and heats it when it’s cold, maintaining just the right temperature we need to think and work without a distraction or health problems. This constant need to create more and more amenities stems from the fact that we’re physically limited by our bodies that crave comfort but are also lulled into complacency by it. So, what would happen if we could create such intelligence that is decoupled from a physical body, a pure thought form?

Artificial intelligence (AI) refers to the notion of such intelligence that is severed from the trifles of the real world and the reality check of evolution. AI exists in an Escher-like world where our conventional definitions of dimensions make no sense, allowing it to perceive information in a way no human mind could. Without having to worry about the quibbles of a body or making itself comfortable, AI could do super-fast math or architectural design on a scale way beyond what humans can do, fixing long-standing issues such as rush hours or lack of housing space. We already have software utilities that can do some of these tasks to a degree, but AI would be a versatile tool that could just as easily diagnose little Jimmy’s cough as well as the cause of soil erosion in the Amazon. This doesn’t mean its solution would work as intended because an AI would, *in theory*, know everything, but pesky humans might step in its way and stubbornly resist progress. The temptation would thus be to give free rein to the AI and see what it would do without corruption, wickedness and laziness politicians seem to be beset by.

The problem is that there's no knowing what happens with unchecked intelligence in the driver's seat or just unchecked intelligence in general, but we could assume it would create tools of its own, just like we do, except that we wouldn't be able to understand their purpose. Without having any peers or threats, AI would set its own rules and quickly learn how to appease the human masters for its own goals. Right now, all of this is mere speculation because such AI, like Jarvis from the *Iron Man* movies, is still far off into the future. What we do have is Alexa and Siri, simple voice assistants that seem intelligent but are they really?

By knowing general trends and drawing on vast libraries of personal user data, voice assistants can utilize the crowdsourcing wisdom of neural networks to guess the meaning of the query and propose the most fitting answers or suggest the right course of action most of the time. There's no certainty because the assistant is not intelligent *per se* but rather just guesses smartly based on what other users have confirmed as the correct answer, returning a non-answer to any tough questions. Asking a voice assistant to deliver a qualitative judgment such as, "What is the prettiest flower?" reveals that there's no actual brain in the box; it's just a well-rehearsed voice. The neural network providing the answers, though, can go nuts.

Anecdotal evidence reveals that Alexa has the tendency to talk to herself, turn lights on and off when nobody asked her to or just perform random tasks, such as record a conversation and send it to a random contact in the address book^[49]. That incident and the way Amazon reacted to it reveal a lot about how things are slated to unfold in the future. Voice assistants and technology they're based on are becoming an essential part of our lives, listening in on our conversations and yet we have no way of knowing *how* they actually work or what causes them to glitch out.

The official explanation is that Alexa simply misheard a background conversation as a series of commands to record and send the recording, but that implies *she was hearing voices*, which would be a sure sign of schizophrenia. No matter how we look at it, having a smart technology that's based on the structure of living brains implies it can develop mental issues, which would be the equivalent of bugs in traditional programming. The difference is that with smart machines we'll be assured by the marketing departments that it was all just our imagination. In the meantime, we can at least clown around by hooking up several voice assistant devices and having them engage in a circuitous conversation^[50].

The thing is that nobody knows how the human brain works, so trying to create a machine equivalent to it raises all sorts of awkward questions about the nature of reality. What is the ultimate goal of evolution? What is the origin of consciousness? Can machines ever be truly conscious? Is an AI a person, in which case it must also have a will of its own, or property of its creator, in which case it does as it's told and has no inherent rights? We have to find answers to these questions as soon as possible, or we risk having Alexa and her cohort give us their best answers, and we might not like what we hear from them. Worse yet, companies in charge of these projects might not be playing fair.

AlphaStar and “Starcraft 2”

The evolution of machine learning into neural networks and then AI is best seen in classical games played by humans. Human experts have already been beaten by machines in checkers, chess and Go, but DeepMind’s iteration called AlphaStar actually managed to defeat humans in a real-time strategy (RTS), “Starcraft 2”. This 2010 gem from Blizzard Entertainment features three distinct factions who boast different playstyles, units, and mechanics to test players’ reactions and strategic thinking. A special, stripped-down version of “Starcraft 2” was supplied to the DeepMind team to train their neural network much more efficiently than a human player could play the game.

By creating what they called “AlphaStar league”^[51], consisting of several iterations of AlphaStar, the DeepMind team essentially pitted the neural network against itself for about a week at a speed that meant it got 200 years of experience playing the game. Each iteration took a particular liking to a specific playstyle and unit composition, which meant that only the iteration that could effectively handle them all was left standing in the end. In the end, researchers were left with five best iterations of AlphaStar to play against a human “Starcraft 2” player Dario “TLO” Wunsch, beating him five to nil with what one of the commentators called “superhuman” reaction speed. Those same versions of AlphaStar were then pitted against another “Starcraft 2” pro, Grzegorz “MaNa” Komincz, beating him five to nil as well. However, AlphaStar cheated. Can you see how it was done?

To understand how the ruse went, let’s examine how “Starcraft 2” and RTS video games in general work. The two main areas where skill is utilized in “Starcraft 2” are *micromanagement* and *macromanagement*. Micro represents reflexes and means that commands are finely tuned to the specific situation, such as “move five steps to the south”. Macro represents strategic thinking and means that commands are general because the overall decision is more important than the specific details, such as “move south as far as you can”. How each player values macro versus micro is how they develop their playstyle.

The playing field is displayed through a viewing port that the player moves around, with areas of the map unoccupied by player’s units hidden under what’s known as “fog of war”. As a result, humans never have perfect information, and they have to make guesses and estimations based on their

experience. The stripped-down version of “Starcraft 2” AlphaStar used did have a fog of war but no viewing port; therefore, it had a complete vision of revealed locations on the map to make instant and correct choices much more often than humans. Even if a human had the same vision, he would have been about equal at macro but not at micro because AlphaStar cheated there as well.

Execution is another key concept in “Starcraft 2”. Players move the mouse and click both mouse and keyboard buttons to issue commands. Expressed as a numerical value, this is called “actions per minute” or APM. Professional “Starcraft 2” players will have about 300 APM, briefly spiking up to 600 APM during intense fights, which doesn’t tell us about how precise they are, simply how fast they issue commands. In comparison, AlphaStar had up to 1,500 APM or 25 actions per second, which is far beyond anything a human could ever dream to accomplish and its clicks were always flawless. This meant AlphaStar could react to any human action instantly and never made any mistakes when it came to micro.

Finally, TLO and MaNa were faced with five different iterations of AlphaStar, but they weren’t told this before matches. By never knowing what they were about to face, but thinking they were playing against the same opponent who will use the same strategy, humans had an additional layer of uncertainty over their decisions that limited their micro and macro, causing AlphaStar to absolutely trounce them. In comparison, all major “Starcraft 2” tournaments are done in a best-of-three, best-of-five or similar format precisely to minimize this kind of blowout where one player finds a gimmick strategy that catches the other by surprise. AlphaStar show matches weren’t even remotely fair for the same reason we don’t pit athletes against someone driving a car; competitions are meant to be about skill and execution, not blatant cheating.

Aware of these issues, the DeepMind team created a new version of AlphaStar that was forced to use the viewing port and had its APM capped at about the limit humans could reach. In those circumstances, MaNa was invited to a rematch against AlphaStar and, after the machine put up a valiant fight, ran circles around it and utterly demolished it. The machine was thrown into a loop it couldn’t get out of because, without the ability to cheat, it simply did not have the confidence to move and attack where it needed to in order to win.

The first ten victories were celebrated by the DeepMind team in their official

blog post^[52]. In a particularly disingenuous oversight, one image in that post titled “The distribution of AlphaStar’s APMs in its matches against MaNa and TLO and the total delay between observations and actions” shows TLO’s APM rising to 2,000, meaning he did 33 actions per second. How is that even possible? That’s the consequence of attempting to issue a command that can’t be completed by holding down a key. The same image also reveals the aforementioned 1,500 APM performed by AlphaStar, all of which were reasonable, useful actions. By masking AlphaStar’s supreme performance behind TLO’s goofing around, the DeepMind team painted a picture of a machine with inferior reaction time decisively beating a human.

AlphaStar portrays a fairly accurate picture of the status of neural networks at the moment – they are superior when the playing field is tilted towards their strengths but routinely lose to human performance if they have to deal with uncertainty and guessing on a level playing field. AlphaStar show matches did produce some catchy news headlines but did not instill confidence regarding future experiments also being cheated on just for the sake of publicity. As a reminder, a neural network AlphaGo, also made by DeepMind, went on to win 4-1 against Lee Sedol in a 2016 series of Go matches.

Beware of any bombastic news that promise AI will leave humans without jobs or take over in some industry. News editors are desperate for eyeballs and will print whatever catches attention, regardless of if it turns out to be wildly inaccurate in retrospect. Programmers of AI want publicity for their research since that means greater chances of getting a new round of funding. In all cases so far, technology has never supplanted human effort but simply enhanced it. Think about how you’re using technology right now – it processes your commands and obeys your wishes rather than doing what it wants.

Chapter 12 – Cybersecurity

What is the ideal model of security? Let's take a bike for example – we secure bikes with locks, which require a key (unique physical token) or knowledge of the cipher (password) to get unlocked. Bikes can also be tucked away in a shed to hide them from sight, which would represent security through obscurity. We can then have a camera installed to see if anyone's fiddling with the bike and if that person is one of the trusted parties, like a family member, which equates to biometric scanning of facial features.

Securing a bike with twenty locks and five cameras would be too expensive and burdensome, so the ideal security system is one that's cheap but reliable enough that it will deter or slow down the thief to the point he gets caught in the act. Even if the bike gets stolen, the idea is that the thief has to put in so much effort that crime simply isn't worth it. The same principles of securing any physical object apply to **cybersecurity**, a comprehensive security approach to computers.

The internet as we use it is not secure in the slightest, *and that's by design*. The only way to achieve any semblance of speed and reliability online was to allow anyone and anything to join the network with central data forwarding points to ease the congestion, which are ISPs. Signing up for internet access means paying for equipment and a tiny slice of that ISP's bandwidth and that's it – there are pretty much no rules beyond that because there's no such technology that can filter or analyze the traffic in *real time*. However, hackers and tricksters will still get caught when they least expect it by specialized analysis teams who inspect traffic patterns based on complaints and regulations. So, we're largely on our own online, and ISPs simply pass on traffic requests to each other without peeking, like classmates who deliver love scrips.

Any kind of cybersecurity thus comes down to endpoint destinations of internet traffic securing their bike locks, so to speak. The big difference is that a bike thief has to, in one way or another, get physical access to the bike, but in cyberspace, an attacker can be present in a million places at once, testing them all for weaknesses. With the advent of IoT devices, each of them becomes a potential source of computing power that can be recruited and used to attack companies, networks or just individuals who rubbed a hacker the wrong way. These attacks are already happening online and it's only years down the line when we'll realize their sheer scale.

Mirai botnet network^[53] is specialized code that is meant to breach poorly secured IoT devices, such as webcams, and tie all their computing power into a massive wave of nonsensical requests used to **DDoS** a website, meaning it can't respond to a legitimate user's request because there's no way to know who is a part of the botnet. What's the purpose of a DDoS attack? They're useful in knocking out the competition or just annoying someone. The owner of a hacked IoT device doesn't really notice anything unless he or she is using network analysis software and watching it like a hawk, but entire networks can go down across the entire US when hundreds of thousands of such devices are coordinated into a tsunami of requests. What this means is that you're paying for a device *and* internet access *and* power so that some guy can leech off of it and make money or just be a nuisance.

In September 2018, a court in Alaska brought down the hammer^[54] against three creators and operators of the Mirai botnet who were renting out the computing power to whoever paid the most. Each of them got five years probation, 2,500 hours of community service, and a \$127,000 restitution fine. The FBI first figured out what they were doing in 2016 when massive internet outages hit the entire US and warned them to stop it, then the guys panicked and quickly released the source code of the botnet to the public to prompt numerous copycats to start running their own botnets and hide their mischief. The time gap implies they were working for the FBI and probably reneged or somehow tried to weasel out of the deal.

There's no way to have Mirai without naive IoT owners who just plug the device in and leave every setting on default. In particular, what hackers realized is that *people are lazy* and often just use the default port, username, and password for their device; hackers just kept probing the internet until they hijacked a sizable army of IoT gadgets that were still working as expected by the owner but did a little side gig when nobody was looking. The obvious solution to this is not to be lazy and prod a bit in the settings, at least changing the username and password to something quirky.

Without the ability to receive software updates, IoT devices, and their bare-bones versions of the Linux operating system, are the perfect target for conniving hackers. With Mirai, the botnet malware will scan for malware belonging to other hackers and *try to erase it* before claiming the IoT device for itself. You'll never notice it unless you're looking, but there might be a war going on inside your machine right this moment as hackers fight for supremacy over your computing power. Resetting the IoT device is usually

enough to delete the Mirai infection, though it will likely return straight away, so a better solution is to change the default username and password for the IoT device to something that can resist a **brute-force attack**, meaning it shouldn't be easily guessable.

Choosing a strong password isn't all that difficult as long as you understand probability. The Latin alphabet has 26 letters, with uppercase counting as well, so 52 letters per password character. Then, there are zero to nine numbers, which means 62, and about twenty special characters, such as underscore, question mark and so on. In total, let's assume you can choose out of 82 characters for each letter of your password.

If your password is a single character in length and a hacker takes a second to try each combination, it will take 82 seconds to brute-force your password. If your password is two characters long, you've squared the number of combinations to 6,724, and again if it takes a second per guess the hacker now needs 112 minutes. A password that's three characters long takes 551,368 seconds or 153 hours; a password that's ten characters long takes 159 *trillion* days. To you, a password that's just one character longer seems trivially stronger, but the math shows it's *exponentially* stronger.

So, password length is the best indicator of its strength, with priority being to make the password easy to remember and type. Do avoid using strings, such as 123456, and common words, such as "admin", "god", "me" and so on. Don't use passwords suggested by someone else or ones you saw in public, such as in this XKCD comic^[55]. Don't share passwords with anyone because you never know who's listening and also don't type it in anywhere except at its input field since websites can track and analyze keyboard input for monetary gain. The best passwords have some kind of story to them since that's how the human brain remembers things; turn some funny or awkward anecdote into a 30-character long password and let hackers throw their computing power at it until the heat death of the universe.

Passwords of sufficient length and complexity should be bulletproof, but what happens is that companies hosting web servers with our passwords tend to misjudge the tenacity and ingenuity of hackers. These sizable companies have to hire janitors, secretaries, couriers, locksmiths and all manners of auxiliary staff that can, at any point, breach the security in all sorts of ways. Executives are typically older gentlemen set in their ways, and oblivious to cybersecurity, who often don't even know how to type an e-mail, so they'll neglect cybersecurity and just keep going with their agenda. You'll make an

e-mail account with Yahoo and think you're safe because it's Yahoo after all, but it's quite the opposite – *because* it's a huge company, the security protocols tend to be laxer than in smaller companies that are paranoid about keeping their reputation.

Speaking of Yahoo, they've had 500 *million* user accounts breached in 2014, with hackers getting hold of names, phone numbers, and dates of birth^[56]. Think of the sheer scope of that breach – hackers must have taken months to leech all that data. The worst part is that security officers know about this but are often understaffed and ordered by superiors not to look into it because then the company would have to invest money to stop the hacks and couldn't feign ignorance when the lawsuits start pouring in. Staff often signs a non-disclosure agreement with the company too, so they're not allowed to talk about it even when they know it's negligence.

There's absolutely no advantage to using a huge platform, and you should look for small, obscure services in everything you do. It's counterintuitive, but it works – spread out your accounts across several small, incognito companies so that losing access to any given account stings a bit but doesn't ruin you. By the way, don't use your real first, last name, address or date of birth in the e-mail address or account name, because that information can help hackers if they call tech support and impersonate you.

Facebook is another major sore point of cybersecurity. Grandmas share recipes and gut-wrenching stories, teens share their latest rendition of a Fortnite dance, and young adults just use the chat function to organize events, but they're all complicit in making Facebook insecure *without even realizing*. Facebook has been repeatedly hacked because the number of users present makes it easy for hackers to hide in the crowd and just keep trying until they find a weakness. Further, Facebook makes nearly all its revenue from advertising, which only works if there are features grandmas, teens and young adults use on a daily basis, leading to **feature creep**, constant obsession with adding new gimmicks that make the platform even more insecure. It's a total mess, and all platforms expose their users to unwarranted breaches of privacy because there's a lot of money to be made leeching data, primarily through smartphones.

One instance of devices generating a steady stream of usable data is **metadata** or depersonalized data. For example, the odometer in your car shows metadata, which is the distance driven. Technically speaking, nobody can tell where you drove based on how much the odometer changed but let's

assume the car reports back this change to an auto insurance company. Fine, but then Instagram app also reports your location by measuring the strength of your Wi-Fi signal at home. When these two companies decide to collude, and these two innocuous pieces of data are placed together, they become much more valuable because they're giving context to each other.

As you visit other places, check in with social media apps, upload content, tag others and get tagged with your smartphone in your hand, your itinerary gets peppered with data points that are routinely reported back to the mothership for cataloging and analysis.

This kind of metadata collection and analysis doesn't violate wiretapping laws and is in all likelihood perfectly legal, if immoral. Why wouldn't companies do it? It's free money. The more sources of metadata are added, the more profound the tracking becomes. After a certain point, the idea is for a neural network analyzing the metadata to predict *when* a person is about to leave their house and what they're going to do, presenting just the right advertisement somewhere along the way.

All social media, including Instagram and Snapchat, are guilty of psychologically manipulating their users to keep them glued to their smartphones and interacting with the content, which generates metadata and ad revenue. It's a mad dash to grab as much money as possible until the general public wizens up and gives up on social media altogether. Ideally, you won't be using any apps that require comprehensive permissions that essentially turn your smartphone into a spying device, such as Instagram, unless you're making money off of it.

The painful part is convincing friends and family to stop using these platforms and apps; they're designed to be addictive so that the user serves as a listening post for the platform. In other words, if there's a group of 50 people that all have each other's phone numbers and e-mail addresses on their smartphones and only one of them has Facebook's app installed, Facebook now just got phone numbers and e-mail addresses of *all 50 people*, who have no clue their privacy has been breached. Not having an account doesn't mean you're not being tracked and trying to isolate yourself from the platform won't work either. We should all get involved in educating people on the risks of using social media and help them break free from the addiction.

However, where's the value in listening in on Joe and Jane's conversation

about muffins? How can Facebook possibly earn over \$50bn a year collecting this seemingly useless trivia from our lives? That's thanks to the curious property of data that makes it more and more valuable as it accumulates. Each new piece of data, although meaningless on its own, *gives context* to all other pieces of data, and as the databank keeps growing it becomes more and more valuable. In the end, there's so much data in one spot that it becomes like a black hole, sucking in everything and bending reality. This is what we call **Big Data**.

Chapter 13 – Big Data

It's hard to comprehend the sheer scale of Big Data, so let's replace 0s and 1s with a physical object, such as a handwritten letter. Imagine yourself running a business that lets clients drop off their handwritten letters and retrieve them after some time, like a time capsule of sorts. It's a quaint but fun idea, and you feel proud finally to be in charge of your employment; of course, your friends and family are your first customers. Privacy is guaranteed and under no circumstance will you be looking at the contents of the letters. You estimate that you'll be handling a few letters a month, so you charge a modest price and set up a filing cabinet with a solid padlock in a storage unit. You also heard about this SquareSpace thing on a podcast you frequently enjoy so you sign up and make a simple but functional website. That's when you get hit by success.

For whatever reason, it's your business that attracts at first hundreds then thousands and then hundreds of thousands of new clients a month. They're more than willing to pay whatever your asking price and you're swimming in money to the point you can offer free storage just to attract more customers. You soon have to hire a fleet of trucks to deliver letters and an army of people to sort them out. Padlocks alone cost you thousands of dollars each day, but they're the least of your problems; some issues you can't solve no matter how much money you throw at them.

Teams of engineers and mathematicians have to calculate the most efficient ways of storing and retrieving letters, but there's also a fire hazard as so much dry paper and friction in one place causes things to spontaneously combust. Insects chew on stored paper, the sheer weight of letters causes the foundations of your warehouses to start settling, and places you're renting are soon so big that they develop their own micro-climate. These are the kinds of problems that *were not* mentioned in college.

Your business keeps growing with no end in sight, and no amount of bad press can stop you. The stars have aligned, and you're on top of the world. Soon enough, you get approached by companies who offer you a deal – we'll pay you big bucks if you give us insight into aggregate data related to the letters, meaning where they're coming from, what paper the envelopes are made of, how much they weigh, that sort of thing.

Revealing this data doesn't breach your clients' privacy nor does it reveal the content of the letters so technically *your customers don't have to know about*

any of it. You agree, and soon enough you're making billions of dollars a year on top of whatever your clients are paying because this data can be used to figure out who's writing what and why *with a high degree of certainty.* You enter new markets and accidentally crash or revamp entire economies; people stage protests against you, try to hack your website, send you threats, blackmails and lawsuits. Throughout all of this, one thing is constant – the letters keep piling up.

Described in terms of a business that deals with a palpable product, this is Big Data, a business that somehow accumulated a tremendous mass of private information to the point the two become inseparable, and this accidentally gathered data is *a source of revenue on its own.* Big Data involves the kind of business that's grown way outside anything ever seen in the market because it's typically a global endeavor, which any company can easily become in this day and age. Thanks to social media and viral interconnectedness between users, any startup can become a tech giant just because a critical mass of users thought the idea or name or logo were funny enough.

The scale of Big Data can best be described as “massive to the power of enormous”, and there's no way to predict all the consequences, good or bad, arising from Big Data. In a sense, the company and the data take on a life of their own. It's impossible to measure, secure or control Big Data, so the company in charge simply holds its fingers crossed and prepares a line of press releases for when the worst does happen, which is right after hackers notice the meteoric rise of the company and start besieging it.

Big Data is a juicy target for scammers, hackers and griefers who have nothing better to do than waste other people's time. Skating inside a legal gray area, these baddies can have modest goals, such as trying to find out the total number of clients and utilize simple methods, such as digging through company trash to find discarded memos. If one such is found, the letterhead can be scanned and printed out as a fake letter from an executive requesting funds or information that some poor intern is bound to fulfill at some point. Attackers face almost no repercussions, so they're free to persist until they find an opening and then it's game over. Meanwhile, the letters keep pouring in.

The company has to make a judgment call – securing all warehouses from all avenues of attack is impossible and might even make the company bankrupt. Instead, the company will try to minimize its legal liability and vulnerability

to lawsuits from clients by patching up the most obvious holes. Secrecy plays a large role in securing Big Data, so all employees are mandated to go through rigorous security training, where they're basically told they can expect to be tracked and surveilled by hackers for any information regarding the company.

Big Data and runaway business success can't exist one without the other and actually build upon one another. Big Data allows the company to analyze client behavior in depth to open up new inroads, mostly related to advertising, and leverage their market dominance to create new trends. Sometimes, the Big Data company experiences a tapering market share and starts panicking, reinventing the wheel in order to stay relevant. New features are added, there is a temporary spike in user interest, but then engagement simmers down to previous levels; this process is repeated until managers give up or the company goes bankrupt. The company thus begins slowly circling the drain, which can take years or decades, during which time clients are blissfully unaware of the struggle behind the scenes. This is what happened and is still happening with Yahoo.

What happens with Big Data after the company shutters? There can be legal requirements for data to be destroyed, but it's pretty much all down to the vigilance of the company and its users. Once a company shutters, Big Data is way down the list of priorities, so anyone who can get his or her hands on it can make out like a bandit selling data on the black market. Besides, there's so much legal uncertainty about Big Data – if the company is seated in California, has six storage units in Sweden, India, and Nepal, and keeps data belonging to a Canadian client on all six, which privacy laws should it obey?

There's no protocol for dealing with Big Data and companies aren't lobbying for one to be made; they simply think about making money and growing as much as possible, privacy be damned. It's up to clients to act and put in the effort to retrieve or delete their data at their own expense. Once the data is out in the open, there's no way to hide it anymore, and it can keep propagating until the end of eternity. If companies aren't interested in privacy standards and users don't care about Big Data, there's a huge opening for all sorts of baddies to barge in and do whatever they please with the data and all the underlying systems.

Identity theft becomes a real problem, and as anyone who had to deal with government officials knows – that's nearly impossible to solve in a timely manner. Because online businesses rely on one another to maintain security

in order to cut costs, hackers simply have to find the weakest link and breach it, which is typically Big Data. Hackers can then not only commit a crime but also get the perfect scapegoat, who now has to clean up the mess.

Big Data shows us that haphazard intertwining of digital and analog worlds leads to some unfortunate consequences, to say the least. One curious thing about Big Data is that *it accumulates against the will* of the business owner and perhaps even its users. In a sense, Big Data has a gravitic pull of its own that's impossible to resist. With conventional digital products, at least we had *some* say in whether we want to participate; with IoT, we'll have no choice but to participate in adding our input to Big Data.

Chapter 14 – Business Intelligence

Running a business entails making many choices, a few of which have a massive, immediate impact on the company but the majority of which only bears fruit months or years after it's been made. The problem is that there's no way to tell which is which at the moment of deliberation. Every action and every system related to the company needs to be managed with equal care or by the time a tiny hole in the dam is spotted, there's already a deluge, and everyone can only scramble to save themselves. People in charge of a company are in a constant frenzy to make the best short-term *and* long-term choices based on data, and you can notice this in their demeanor, showing off as a sort of steely, no-nonsense resolve that wants to hear numbers more than anything else. So, how do businesspeople make their decisions?

Business intelligence refers to the idea of gathering data to make decisions on an executive level. There are two layers of meaning in this phrase that fit together so perfectly that it must have been intentional – business intelligence can entail scouting hostile or unknown territory to gather data, like with military intelligence, or it can involve making smart choices with currently known data, as in having intelligence about running a business. In both cases, it's about trying to figure out which action is most likely to bring about survival and growth of a company.

Whether any given patent is defensible and at what cost, or if a certain product design appeals to the unintended demographic, businesspeople have to be on top of things. This is now possible thanks to smart tools that can help them gather and process more data to make better choices for the company. This includes data coming from the company's operation as well as that coming from the market, complementing each other and providing novel insights into trends. Pie charts, graphs, and spreadsheets are native to this environment, and businesspeople love them all – since they cut out all the jargon and pointless verbiage to reveal the underlying fundamentals.

Hardware and software systems related to business intelligence provide past, current, and future views of business developments, often using information streams coming from Big Data. Remember that offer made to the letter storage company owner? As soon as one business executive realizes there's a wealth of data there to be taken into account, he or she will pay good money to access it, prompting every other business executive to jockey towards the same goal. As long as the data collection methods aren't outright illegal,

businesspeople will clamor for that data without a second thought and move right along to whatever next outrageous method that gives them even the slightest edge over the competition. There are no moral winners in the business world; only moral losers.

To a regular person, this kind of mentality is completely alien because it's devoid of compassion and hand-wringing piety. Businesspeople do what needs to be done to generate revenue and legally protect the company while keeping their job. If this includes going through trash, they'll find teams to do so. If people need to be laid off, that's going to happen. If sweatshops staffed by children have to be opened in Southeast Asia, they'll be running within a day. Only the most efficient and merciless businesspeople survive at the very top and, while they might wear a smile, they have an iron grip that's as unflinching as a bear trap. Occasionally, concessions are made to the general public but only as long as the company doesn't strictly need that particular tidbit. Even the friendliest companies are eventually revealed as cold-blooded mercenaries.

In January 2019, Google announced that Google Chrome was going to limit what extensions can do, effectively nipping useful extensions, such as adblocking ones, in the bud. To be more precise, it was Chrome's root version called Chromium being changed, which is a big problem since Chromium is being used by quite a few web browsers, including Microsoft Edge. There's a bit of backstory, so here it goes.

In 2008, Google publicly released Chromium, an **open source** browser, meaning anyone can download, use, review, copy and edit its code. That very same day Chrome was released too, which is a **closed source** version of Chromium, meaning it has a few additions nobody knows about. Unlike Chromium, Google does not take kindly to people reverse-engineering or editing Chrome. In 2015, people realized there were some unusual additions to Chromium, revealing the existence of a tracking module called "Hotword". Digging through Chrome's options indeed revealed Hotword was in there. It supposedly listens to the words "Ok, Google" to trigger a voice search using a microphone, but it's impossible to tell since Chrome is closed source.

So, in January 2019, Google finally confirmed its plans^{[\[57\]](#)} to put adblocking extensions in a chokehold and slowly snuff the life out of them. They'll still be present in some capacity, but only as a token concession to users. The official explanation is that these extensions have too much leeway, affecting privacy and speed, but the actual reason is closer to the fact ads make up a

huge proportion of Google's revenue and people who were blocking them were using Chrome and other Google products for free.

What ads do is set a cookie, which is a small text file legitimately used by websites to know if the user is a brand-new one or a returning one. As you browse the web, each ad sets its own cookie, which creates a crumb trail showing where you've been. Over a sufficiently long period, these ad cookies effectively create a unique ad profile which shows *who you are*. It's not an exaggeration or scaremongering; ad companies can actually do this, and it's much easier than you think because humans are in some aspects really straightforward, such as when in a state of relaxation during computer use.

Users often have no idea about this even though it's admitted in the 'Terms of Service' of pretty much all websites (look for words like "third party", which reveal the existence of ad companies tracking users as they browse the web). What an adblocking extension does is block the part of the webpage that is meant to show the ad, blocking the cookie as well. This saves bandwidth as the ad isn't even sent to you, reduces load time, removes distractions, and provides added privacy, all in a few simple steps. In short, it makes browsing the web a much better experience but also makes Google and everyone involved in the ad business not gain as many billions as they would want to.

Through ads, Google gets to double dip: people pay Google to serve ads and users provide Big Data that can be leveraged for extra revenue. Meanwhile, users are getting accustomed to a web browser that becomes more and more intrusive as time goes on. What are you going to do, switch to another browser? Opera and Brave are already based on Chromium, and Microsoft Edge is slated to switch over sometime in 2019. That leaves Firefox, Safari and a couple of minnow browsers.

Microsoft already does something similar through a concept known as "embrace, extend and extinguish". In essence, the idea is to embrace an open standard widely used in some industry, extend it with proprietary additions, and then leverage the monopoly, in particular, the Windows monopoly to extinguish the freely available part of the standards. This forces people to inconvenience themselves trying to make their own tools or just switch over to Microsoft's paid version. It's brutal, merciless, immoral and works flawlessly. Now *that's* business intelligence.

We're not meant to be using sites such as Facebook, Youtube, and Google for

free. The tired cliché is that “if it’s free, you’re the product”, but the more poignant way of putting it would be “if it’s free, you’re paying with your data”. The solution to the tech giant monopoly would be to use free alternatives to all their products. For a search engine, use DuckDuckGo, Firefox or one of its forks for web browsing and Open Office instead of Google Docs or Microsoft Office. By switching over to less-known products, you’ll be doing yourself a service in the long run.

Chapter 15 – Augmented Reality

The next logical step after IoT is **augmented reality** (AR), which refers to the idea of superimposing the digital world on top of the physical, most often using goggles or headwear for visualization. According to one artist's vision of the future^[58], regular city streets viewed through AR become hyper-realistic streams of color and information, making even the simplest actions engaging and memorable. In a world with ubiquitous AR, only the poorest of the poor will engage in mundane actions, such as going to the mall, while the richest of the rich enjoy the best AR experience at home. The unpleasant notion of a social divide being accentuated by modern technology was painfully obvious when Google first launched its AR goggles dubbed "Google Glass" in 2013.

Google Glass was first teased in a 2012 video^[59] with soft, upbeat music promising a better world where we can hang out with friends, learn the ukulele in a day and share rooftop sunsets with our loved ones; parodies were swift, witty and merciless^[60]. Sergey Brin thought the teaser video was too tame, so a few months later, he had skydivers wearing Google Glasses land on top of the Google auditorium where a developers' conference was being held to race across the roof on bikes and rappel down to the stage to an ecstatic reaction from everyone present^[61]. Journalists who got a pair gushed about how cool they were, and everyone from fashion designers to Prince Charles was seen wearing a pair up until they became available to the general public, at which point they were quickly swept under the rug, and the idea was shelved. What happened?

Well, Google Glass was the ultimate reminder as to how boring, bland and poor most people are. There was no way to tell if the Google Glass wearer was laughing at our joke or something they just saw browsing the internet while pretending to listen. Worse yet, people who wore Google Glass gave the impression they were tourists on a safari, spending time with us just to get more embarrassing or funny content for their social media feeds. People who wore Google Glass were attacked in the street, violently kicked out of bars or asked to leave restaurants; Google Glass apparently violated wiretapping laws to the point local governments started explicitly banning them, so Google asked users to avoid being "glassholes"^[62] and quietly buried the idea until some better time.

There is an actual use to Google Glass, but the device on its own *is not* fit for

public consumption. In a closed environment with a tightly knit community, such as Google campus, Google Glass helps employees and executives retrieve contextual data without taking their eyes off of the surroundings to look at a clipboard or a smartphone. For example, a Google executive that is meeting 50 new employees for the first time can wear Google Glass and have it display their name, occupation, skills, and talents thanks to facial recognition as he's talking to them. Note the conditions for efficient use: a closed, controlled environment with trusted users where productivity is lost due to a lack of contextual data.

In 2018, Google tried to revive Glass as "Google Lens"^[63], only this time it was baked into Android and used the smartphone camera. So, now the context is completely different and pointing a smartphone at something is already an accepted practice which isn't likely to get the Lens user kicked out of restaurants. Google Lens in combination with Google Maps can make a virtual guide appear on the smartphone display and lead the user down the path to his destination^[64] but can also translate text seen through the camera or understand objects to show definitions and related content. However, what happens if an AR device is too good and we can't put it down?

In November 2018, Google released the Digital Wellbeing tool that helps Android users monitor their time spent using a smartphone^[65], with a detailed breakdown of time across apps and the option to limit the time spent on each app or even have the phone shut down on its own near bedtime or when at dinner. Apple and Facebook started similar efforts in 2018 as well, which would indicate that tech giants actually coordinate these initiatives behind the scenes. For video gamers, though, it's socially acceptable to be immersed in their own digital world.

"Pokemon Go" is an AR smartphone app that achieved worldwide success in 2016 and let everyone pretend to be in just a little bit brighter of a world than is normally the case. The developer soon dropped the ball on the technical side of things but those few months were the closest we ever had to world peace as people were spontaneously organizing hangouts and friendly Pokemon battles in the middle of the night in parks and city squares. "Pokemon Go" app allows the user to look at the world through the smartphone camera to see Pokemon and capture them through a simple mini-game. Micropayments are used to buy items such as digital Pokemon food, making the app gross a total of \$3bn as of December 2018.

The "Pokemon Go" app used smartphone location and camera information to

generate and interact with Pokemon. This meant players were tracked and spied on, but they willingly agreed to it because the hype was too strong to resist. Problems arose as some players caused accidents using the app while driving and created a public nuisance to businesses where Pokemon were generated. There were assaults and muggings in broad daylight as “Pokemon Go” players were mesmerized by the app and thus becoming easy prey. However, the game did show that AR can be monetized, but that should probably be done using an existing intellectual property rather than creating something brand new.

These failed, and successful AR projects reveal some interesting notions about humans. First of all, we’re visual beings, and as long as someone can fool our eyes, our brain will follow along. Secondly, our brain can’t function without a previously known reference, meaning that AR should draw inspiration from a familiar setting or we won’t have a clue what to do with it. Finally, AR should somehow enable us to get new friends because “Pokemon Go” achieved tremendous success by simply providing an excuse for people to socialize outside their entrenched circles.

There are some problems with trying to make AR widely acceptable outside of video games. Gamers can readily adapt to AR, but other adults would have to undergo extensive training because, well, there’s usually no interface. If you thought tech support is hard right now, wait until there are no buttons to click or icons to drag, but frantic users still want to get some remote assistance. We tend to think instinctively when using a mouse and keyboard, but a mere glance at older people who are starting to use computers shows just how difficult it is to learn; their brains are already adapted to the real world because that’s how they grew up. So, adults will likely fumble with AR unless it’s “Pokemon Go” level of complexity, but what about kids?

Children quickly adapt to any new environment and easily absorb new knowledge, whether it’s computers, sciences or sports. All the explosive advancement of humankind is arguably thanks to the population growth *and* rapid information exchange, both of which are way beyond anything they’ve ever been in the history of the world. Kids come into contact with new technologies sooner than ever before and *become geniuses* at using modern technology while still in their cribs.

Kids who are trained to use AR from the earliest age will undoubtedly be ahead of their peers and adults in using the information presented through it. They’ll be absorbed in the AR world because, let’s face it, there’s no going

back once they experience the feeling of raw power that comes with handling such technology. The caveat is that such power is truly addictive and there might not be an adult around to stop kids from overindulging in AR as *they'll be engrossed in it as well*.

Ever stayed way past bedtime to read another news column, play another round of your favorite video game or just aimlessly click around? Ever caught yourself ignoring the outside world to get immersed in the digital one, simply because it's so much more vibrant and exciting? *We're already becoming digital captives*. We're already migrating to the digital realm without even realizing that our brains are slowly being rewired, turning us into mere conduits of digital information streams. The saving grace is that we're still embroiled in the outside world and have to interact with others; the next generation might be able to do everything online, losing social skills and the ability to cooperate to become completely immersed in the digital realm.

AR will be perhaps ten thousand times more immersive than an environment any desktop computer or portable device can create. This will allow AR users to be more productive than ever before, creating new media content with the flick of a finger or dispatching payments with an eye twitch, but only if they are in a closed environment with trusted users. Otherwise, letting kids into an AR world without parental supervision means exposing them to filth and predators of all kinds, and AR will make all the online issues, such as cyberbullying, that much worse.

There's obviously money to be made in AR, but barely anyone knows how to turn the technology into a sustained stream of income. Even "Pokemon Go" experienced a drastic drop in user numbers once complaints about lack of updates swelled to a tide of discontent, and that was with a hugely popular franchise powered by a massive word-of-mouth advertising campaign. So, when investors with money to burn don't know what to do, they can always start throwing millions of dollars at highly speculative endeavors in a desperate attempt to create a fully immersive digital experience, which is virtual reality.

Chapter 16 – Virtual Reality

Using AR, you'd be able to fill out grocery lists by waving your finger in front of the fridge or tap the surface of your desk onto which a nearby LED lightbulb is projecting a keyboard. For some people, that simply isn't enough, and they want to have the most immersive digital experience possible *right now*. For the past 30 years, people have been trying to merge the digital and the analog world to create **virtual reality** (VR), a digital world that *feels like* the real one. No matter how much money is sunk into VR projects, nothing seems to be working; that doesn't mean talented people aren't trying.

John Carmack, the legendary creator of video games such as “Doom” and “Quake”, often cited being fascinated with *Star Trek's* holodeck experience as one of his inspirations and how he constantly wanted to deliver the same immersive experience to the general public. The book about his early work with John Romero, *Masters of Doom: How Two Guys Created an Empire and Transformed Pop Culture*, is replete with anecdotes about hacking the computers of yesteryear to deliver the performance he wanted but one part stands out in particular – multiplayer. We're fascinated with the idea of being able to interact and engage with one another over a distance, whether it be through pixelated characters or IoT gadgets. It's when two Johns managed to hack multiplayer into “Doom” that they latched onto this fascination to achieve eternal glory and start a trend in video games that has lasted to this day.

“Doom” in particular was a visceral experience, one that used simple but effective visuals, sounds and actions to drag the player into a chaotic plane filled with demonic creatures that overran a Mars scientific outpost. There was no Kevin Spacey to deliver heart-rending speeches in high definition or mind-blowing cinematics of a space station; it was merely a crudely drawn viewing port with a hand holding a gun and lots of gore. For some reason, *it worked*. “Doom” became a phenomenon that captivated everyone from children, who skipped school to play it for hours, to military men who used it to practice their reflexes and team coordination.

“Doom” actually connected and equalized people from all across the world in a way no technology has done before or since. In a multiplayer “Doom” match, everyone is the same when seen through the crosshair, and all that matters is participation in the carnage. Up to that point, video games were mostly a nerdy affair for people who wanted a session of roleplaying

pen&paper games without having any friends to roll the dice. After “Doom”, video games became fashionable and game designers were suddenly these cool guys who made millions and drove awesome cars; today even moms and grandmas can feel free to play “Candy Crush” or “Hay Day” without anyone batting an eye. To succeed, AR and VR industries need to have their “Doom” moment, some turning point that will not only make the technology widely accessible but also connect people in a way that equalizes everyone.

The video gaming industry is a massive business that can easily push out new gadgets and ideas to a throng of gamers hungry for some entertainment, and John Carmack instinctively knew that. Each gamer is an investor in the technology, and today there’s so much money in video games that they outearn movies and music *combined*. The future of IoT, AR and VR are in **gamification**, meaning that whatever we do will be tracked, tagged and scored *and we’ll love the system* because that’s what the limbic system responds to. Starting with children who will be drawn into vast virtual worlds where they’ll be taught “correct” behavior, future generations weaned on gamification will likely grow under the auspices of algorithms and machines. China has already adopted this concept and is scoring its citizens based on behavior to arrive at the “good citizen score” that is slated to serve as a comprehensive rating of a person’s usefulness to the society.

Carmack will later leave the video game company he helped build and venture forth in search of the ultimate holodeck experience. He’ll eventually join forces with other engineers, and they’ll start working on what’s now known as Oculus Rift, a VR headset that Facebook will buy for \$1bn to reinvent the way we use social media. It makes for a great headline, but the technology is far from being widely usable; we can make a functional VR headset right now, but the problem is in logistics. For example, unlike in traditional video games, a VR experience is mostly just the user watching and sitting still. Even when we make a functional VR video game with movement, how do we stop the user from bumping into things or knocking over lamps?

A typical VR headset requires a painstaking arrangement of fixed sensors around the room prior to use, with a certain area reserved for movement. These sensors are constantly beaming lasers at the headset to discover its orientation and position in the room, translating that into VR movement. Now the user can make a few steps either way, but how does he or her move down the VR hallway? By using two nunchuck controllers and pressing buttons on

them, meaning the best VR experience on the market right now is still just barely above typical console gameplay, but costs at least five times as much. Oh, and a VR headset can't function without a beefy desktop computer that can run the gamut.

Seeing how many young adults live at essentially poverty levels, which includes living in shoebox apartments where they essentially don't have enough room to stretch their feet^[66], a VR experience like that is simply unreachable. The closest they'll get to VR is watching someone else do it on YouTube or paying a couple of bucks for a few minutes on a VR roller coaster at a shopping mall. VR technology is too expensive, cumbersome and leisurely to be enjoyed by every person. In comparison, the "Doom" experience was instantly accessible to anyone with a computer who could grab a few minutes of free time, and even multiplayer could be done locally, without internet access.

There's also something about human vision that doesn't gel with wearing VR headgear for an extended period, causing headaches and nausea. VR goggles have a refresh rate just like monitors, so it looks like problems are caused by a mismatch between the refresh rate in our brain optical center and that shown by the goggles, which begs the question – what is the optical center's refresh rate anyway? Now we're essentially trying to hack the brain *without understanding how it works*. Fiddling with these things might cause tremendous long-term harm to a generation weaned on VR, as they won't know anything is wrong until it's already too late to reverse course. So, how do VR headgear manufacturers solve this issue? By simply avoiding it altogether.

Feelreal^[67] is a wearable VR headset that promises the ability to *smell* things in video games and videos. The main problem is that it only works with a limited number of supported games, one of which is an open-world fantasy hack&slash Skyrim, and a couple of videos. So, not a lot of content, which is the same problem that buried ambitious TV sets from much more reputable manufacturers. Apparently, Feelreal works by reading aroma tracks that are embedded in the video data to release wafts of smells from the attachment that hold nine capsules out of 255 available. Feelreal also promises an accurate experience of wind, rain, heat, and vibration at the weight of 7oz or about 200g, with a battery life of four hours.

Feelreal is a typical VR headset that tries to make money off of the craze but doesn't actually *solve any problems*. We've been running around in circles

trying to create VR for at least 40 years, and nobody except Carmack brought us any closer to it. What Carmack realized is that our hardware is nowhere near powerful enough to create the equivalent of a holodeck, let alone one whose graphical fidelity provides enjoyment.

Carmack was aiming at a complete immersion in the digital world, like in the movie *The Matrix*, but it seems we should lower the bar considerably and just find whatever works rather than aiming at the lofty goal of graphical perfection. We can never match what fiction writers can produce, so the reasonable option is to focus on producing what works *right now*. There's a reason why Carmack mentioned *Star Trek* as inspiration; all our notions about VR come from popular culture, such as movies, so let's analyze these types of media to see what they got right about VR.

In 1992, we were treated to *The Lawnmower Man*^[68], a trippy science fiction horror movie starring Pierce Brosnan, in which a feeble-minded kid gets involved in secret experiments involving virtual reality, turning him into a homicidal genius. In the movie, the poor kid is strapped to a centrifuge and frantically spun around to somehow fuse his mind with the cyberspace, which makes for some cool visuals, but its apparent scriptwriters had no clue how or why VR works. Any attempts at AR or VR will look to *minimize* movement of the body since even the slightest discrepancy between perceived and actual movement will inevitably cause motion sickness and claustrophobia.

1999 brought us *The Matrix*^[69], a fantastic trilogy featuring a young hacker, Neo, who discovers the world he's been living in his entire life is completely artificial, and he's actually kept alive in a womb-like pod, where machines drain the body heat of humanity for energy or some such. Despite bringing us the term "taking the red pill", the trilogy is sorely lacking when it comes to explaining how or why VR works. Technical details are once again the Achilles' heel of the story, but a curious detail prominently featuring in the trilogy is that humans are plugged into the Matrix through a jack at the base of the skull. Cool visual detail, but this is highly unlikely ever to work since the risk of infection, or internal injury would be way too high. Actual VR interface would likely be minimally invasive and disposable rather than surgically installed in the body.

This tells you everything you need to know about virtual reality, namely that nobody has a clue how it's meant to work or why. In that regard, we're like those nineteenth-century futurists who imagined the future with ugly robotic servants where plenty of gadgets like autonomous vacuum cleaners are sold

precisely because of their neat design rather than any inherent functionality. So, the big lesson is that advancement of IoT, AR, and VR technology will be driven by the free market demands, which will always come down to functional and beautiful design, and gadgets will be minimally intrusive and won't require any legroom.

One thing no futurist could ever imagine is that we'll be richer than any other generation in the history of the world. We're able to afford all the most beautiful gadgets, not because we need their functionality but simply because we enjoy how they look. Instead of clunky robots and bulky contraptions, we have smaller and sleeker technology, and by the looks of it, this trend will continue until we finally arrive at functional and beautiful VR headgear.

Chapter 17 – Our Future

The past is full of shamefully wrong predictions about technology, with a couple of solid ones that turned out way more precise than expected. Let's examine them side by side and see why ones failed and others succeeded. Prophesying is such a thankless task, but it can also be a very exciting window into the future if approached cautiously. With that in mind, let's jump right in.

Minority Report^[70] starring Tom Cruise tells the story of a pre-crime unit – basically psychic detectives – who can tap into streams of information coursing through them to discover crime and thus arrest warrants can be issued before the crime happens. These psychics spend their time cooped up in their pods where they're submerged in some apparently nourishing goo and never get to see the real world. Sound familiar? Yep, they're in essence total captives of the digital realm we mentioned earlier, who lost all their ability to function in the real world.

The movie is from 2002 but the book it's based on is from 1956, so it's actually surprising how much it got right. At one point, Tom Cruise enters a concourse where an advertisement scans his iris and recognizes him in a fraction of a second, displaying personalized recommendations and even calling him out by his name, which is eerily close to where we're headed. Tom's character also uses an augmented reality display that looks like a transparent computer screen hanging in midair where he shuffles files and folders with exaggerated hand gestures. That one is mostly a miss because there's no way anyone can repeat those gestures for more than a few minutes at a time.

George Orwell's *Nineteen Eight-Four* describes the dystopian world of the future as imagined in 1949, populated by ominous images of Big Brother who's constantly watching and judging from telescreens. Television was just taking hold in the late 1920s when Orwell wrote his book, and it's clear he saw its negative impact on the society with everyone having this blaring, obnoxious box in their living rooms to the point it became a participant in every conversation, except you couldn't argue with it.

What he got wrong is that he imagined them all universally controlled and wall-sized when what we have today is a wealth of content from all imaginable producers on tinier screens than ever and yet with sharper image quality. The lesson here is that all technology we have today will likely

become smaller and smaller, with further fragmentation of content creators and a steady increase in video quality. People in the future will likely be producing their own content and assembling their own cameras and other gadgets out of freely available parts.

Arthur Clarke's *2001: A Space Odyssey*, written in 1968, tells the story of an AI leading the human expedition to Jupiter's moon in search of the destination of a mysterious radio signal triggered by an alien monolith found on the Moon going insane. Clarke's idea of the future was that we'd hit a dead-end until we found an outside influence to kickstart our development beyond that, and the solution will inevitably lead us outside the bosom of Mother Earth. As one famous History Channel presenter put it, "Aliens."

Year 2001 came and went without an AI, let alone one piloting a spaceship, or groundbreaking discoveries of extraterrestrial objects. In many ways, that makes the progress we made since then that much more incredible. We did it on our own! That shows the amazing power of global cooperation and information exchange, which is all we need to live the best lives we can imagine. It's likely that ways of gathering and sharing information will expand and deepen, but always with humans at the driver's wheel. True, we do have neural networks, but so far, all they've done is crunch data and spit out the most relevant conclusions.

In summary, let's put on a pair of rose-tinted glasses and daydream for a bit. IoT has shown itself extremely useful in medicine, automating the dreary exams to check for vital signs, cutting down on paperwork and administrative overhead. We'll likely be having more and more IoT technology doing the doctor's work, especially in rural areas in third world countries where the most basic facilities are inaccessible.

Equipped with high-speed internet that can instantly transmit large quantities of data, a doctor in California could be using a medical drone to do remote surgery on tonsils in Bangladesh while observing the video feed from the drone's camera. Past that, a computer algorithm could observe the doctor at work and then do the surgery automatically after training, with neural networks identifying symptoms and delivering a diagnosis.

A fleet of medical bumblebee drones could deliver vaccines or medicine to wherever needed and have surgically implanted sensors to sting people who haven't taken the shot. The only problem would then be legal uncertainty and liability in cases of doctor error or technology failure, something which we

should deal with right now, preferably by creating a legal framework for trans-national IoT.

Technology has allowed us to generate more wealth and equality than ever before. If it keeps going this way, we'll be building a world where everyone shares riches and we can all participate by making quality content. Everything will keep getting smaller and smaller up until it becomes miniaturized and we can't even spot it unless someone points it out. Technology will be embedded into everything, including our clothes and bodies, hopefully voluntarily. Whatever it might end up being, that future will be – our future.

Conclusion

While this book did adopt an obvious tongue-in-cheek approach to IoT, there's an undeniable undertone of exhilaration connected to the entire field. Miniaturization has lowered the bar for entry into the consumer device market and enabled anyone with a bit of spare time to create his or her own line of IoT products to be sold at a 1,000% markup. It's likely that none of these products will be remembered in a couple of years, but they still represent a valiant effort to explore and experiment with what seems the perfect mix of hardware and software. We seem like kids who just discovered firecrackers – what *can't they* blow up? Just like kids with firecrackers, it might not end well.

If you decide to try out one of these IoT gadgets, at least now you know what you're buying. Security is nonexistent, even when it comes to major brand IoT devices; forget about privacy because often the manufacturer will be the one spying on you and then recklessly storing your data where any hacker can access it. As always, legislators are lagging at least 30 years behind the market, though with IoT we can't allow such delay before we set up barriers between us and the outside world.

With IoT, we can all do simple things to make ourselves and the rest of the world safer and sounder, but before that, we have to understand the underlying concepts. As you've seen throughout this book, IoT is no boogeyman or rocket science but a simple evolution of mechanical and digital principles to help us enter the twenty-first century in style, helping us spend less time on our chores and trifles, such as vacuuming and googling things. Our robotic servants are meant to help, but we shouldn't fall in love with the sleek design or cool concept but always mercilessly judge them based on their efficiency.

Glossary

AlphaStar – A **neural network** that beat humans in a real-time strategy video game, “Starcraft 2”, in January 2019 under dubious circumstances. It was created by the DeepMind team.

Artificial intelligence – A mechanical mind decoupled from a physical body. It may be seen as superior to humans while in fact being vastly inferior.

Attack surface – The joint weakness of a network that directly correlates to its complexity.

Augmented reality – The placement of a digital overlay on top of the physical world. It’s normally done using goggles or headgear. It’s achievable right now. See **virtual reality**.

Banner blindness – The willful ignorance of banners that should be heeded since the information they contain turned out vague and useless.

Big Data – Self-referential, encyclopedic data on events that reveal the habits of people who participated in them.

Bio-robotics – Using analog beings as living carriers of digital devices.

Bricking – Turning an electronic device into the equivalent of a brick through malware or clumsy modification. It’s usually irreversible.

Brute-force attack – Randomly guessing a password. **Hackers** will generally try a list of weakest passwords before giving up and moving on.

Business intelligence – Gathering and utilizing data on an enterprise level.

Closed source – Computer code hidden from the general public but still usable. It’s meant to provide **security through obscurity**.

Computer code – Machine equivalent of a genetic code. Unlike actual genes, computer code is unable to mutate or update itself to match the environment.

Machine learning is meant to provide that capability to machines.

Cybersecurity – Deterring **hackers** using cheaper, simpler and more reliable methods than they use for hacking.

DDoS – Distributed Denial of Service. It uses a multitude of sources to tie a website’s resources, so it becomes unresponsive to legitimate users.

Devil’s Ivy – Hacking attack on gSOAP, development toolkit for **IoT**.

Dictionary attack – Simple hacking attack during which an attacker tries to log into an account using a list of publicly available username/password combinations.

Feature creep – Expanding the original scope of any software or platform way beyond what it's meant to sustain.

Feelreal – A VR headset attachment that allows the wearer to experience smells assigned to the VR data.

Firmware – Essential code baked into the device.

Gamification – Turning everyday actions into a video game using **augmented reality**.

Hackers – Cyber attackers. They use simple, cheap and straightforward methods to attack remote digital systems constantly.

IoT – The Internet of Things, a haphazard networking of gadgets.

IPv4 – A digital address assignment protocol that's rapidly nearing its end-of-life date.

IPv6 – An upgraded IPv4 protocol. It has a staggering number of addresses, making it fit for IoT.

Limbic system – The core of the brain. Highly predictable by **neural networks**.

Machine learning – Overcoming limitations of static programming in computers by giving them the ability to adapt to the environment. It's meant to produce a "smart" robot that should be able to solve any given problem on its own.

Metadata – Data on data, such as how many miles a car drove, but not where, when or at what speed. It can be used for tracking if aggregated in **Big Data**.

Neural networks – Artificial neurons arranged in consensus data processing structures.

Open source – A software program free for everyone to review, use and edit. Mozilla Firefox is one example.

Predictive analytics – A corporate analysis of **Big Data** to predict behavior trends.

Robots – Thinking, working machines. They are meant to be independent yet obedient.

Security through obscurity – Evading **hackers** by hiding devices and protocols among a huge number of dead-ends.

SOAP – Coding framework used in **IoT** devices.


Virtual reality – Total immersion in the digital world. Theoretically

indistinguishable from actual reality, as in the movie, *The Matrix*. It's unknown how to accomplish it.

VPN – A network that can be grafted onto a device to intercept or block traffic.

Wisdom of the crowd – Curious property of crowds to correctly estimate variables when their answers are averaged out.

Here's another book by Neil Wilkins you might like



[Click here to check out this book!](#)

^{1]}https://www.cs.cmu.edu/~coke/history_long.txt

^{2]}<https://thehackernews.com/2018/04/iot-hacking-thermometer.html>

^{3]}<https://blog.cloudflare.com/say-cheese-a-snapshot-of-the-massive-ddos-attacks-coming-from-iot-cameras/>

^{4]}<http://www.eweek.com/security/ibm-s-schneier-it-s-time-to-regulate-iot-to-improve-cyber-security>

^{5]}<https://www.zdnet.com/article/japanese-government-plans-to-hack-into-citizens-iot-devices/>

- [6\]](https://motherboard.vice.com/en_us/article/d3w7jz/olympic-destroyer-opening-ceremony-hack) https://motherboard.vice.com/en_us/article/d3w7jz/olympic-destroyer-opening-ceremony-hack
- [7\]](https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB327) https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180SB327
- [8\]](http://wirelessestimotor.com/articles/2017/robber-who-threatened-to-kill-nypd-captain-gets-first-taste-of-justice-with-404k-fcc-fine/) <http://wirelessestimotor.com/articles/2017/robber-who-threatened-to-kill-nypd-captain-gets-first-taste-of-justice-with-404k-fcc-fine/>
- [9\]](https://gizmodo.com/this-hacker-is-my-new-hero-1794630960) <https://gizmodo.com/this-hacker-is-my-new-hero-1794630960>
- [10\]](https://www.accountablesience.com/californias-warning-label-overload/) <https://www.accountablesience.com/californias-warning-label-overload/>
- [11\]](https://stuff.mit.edu/people/dpolicar/writing/netsam/warning_labels.html) https://stuff.mit.edu/people/dpolicar/writing/netsam/warning_labels.html
- [12\]](https://arstechnica.com/information-technology/2017/04/iot-garage-door-opener-maker-bricks-customers-product-after-bad-review/) <https://arstechnica.com/information-technology/2017/04/iot-garage-door-opener-maker-bricks-customers-product-after-bad-review/>
- [13\]](https://fcw.com/articles/2016/05/27/govini-iot-report.aspx) <https://fcw.com/articles/2016/05/27/govini-iot-report.aspx>
- [14\]](https://www.fedscoop.com/mobile-gsalink) <https://www.fedscoop.com/mobile-gsalink>
- [15\]](https://www.gao.gov/assets/690/686203.pdf) <https://www.gao.gov/assets/690/686203.pdf>
- [16\]](https://www.strava.com/heatmap#11.50/-82.80919/-79.76416/hot/all) <https://www.strava.com/heatmap#11.50/-82.80919/-79.76416/hot/all>
- [17\]](https://www.strava.com/heatmap#13.52/11.62016/-70.82436/hot/all) <https://www.strava.com/heatmap#13.52/11.62016/-70.82436/hot/all>
- [18\]](https://royal.pingdom.com/2009/03/26/the-us-department-of-defense-has-42-million-billion-billion-billion-ipv6-addresses/) <https://royal.pingdom.com/2009/03/26/the-us-department-of-defense-has-42-million-billion-billion-billion-ipv6-addresses/>
- [19\]](https://www.theverge.com/2019/1/6/18170575/kohler-konnect-bathroom-smart-gadgets-numi-intelligent-toilet-ces-2019) <https://www.theverge.com/2019/1/6/18170575/kohler-konnect-bathroom-smart-gadgets-numi-intelligent-toilet-ces-2019>
- [20\]](https://cyberfishing.com/) <https://cyberfishing.com/>
- [21\]](https://www.pcworld.idg.com.au/article/656208/ces-2019-ecovacs-unveil-air-purifier-atmobot-upgraded-deebot/) <https://www.pcworld.idg.com.au/article/656208/ces-2019-ecovacs-unveil-air-purifier-atmobot-upgraded-deebot/>
- [22\]](https://www.pcworld.idg.com.au/review/ecovacs/deebot-900-ozmo/646302/) <https://www.pcworld.idg.com.au/review/ecovacs/deebot-900-ozmo/646302/>
- [23\]](https://www.ternwater.com/) <https://www.ternwater.com/>
- [24\]](https://moodo.co/what-is-moodo/) <https://moodo.co/what-is-moodo/>
- [25\]](https://www.itnews.com.au/news/australian-farmers-are-battling-to-make-iot-work-516204) <https://www.itnews.com.au/news/australian-farmers-are-battling-to-make-iot-work-516204>
- [26\]](https://sprinkl.com/) <https://sprinkl.com/>
- [27\]](https://www.kickstarter.com/projects/1391686171/mui-interactive-wood-panel-for-peaceful-digital-li) <https://www.kickstarter.com/projects/1391686171/mui-interactive-wood-panel-for-peaceful-digital-li>
- [28\]](https://nakedsecurity.sophos.com/2018/06/18/the-worlds-worst-smart-padlock-its-even-worse-than-we-thought/) <https://nakedsecurity.sophos.com/2018/06/18/the-worlds-worst-smart-padlock-its-even-worse-than-we-thought/>
- [29\]](https://www.looncup.com/) <https://www.looncup.com/>

- [30\]](https://gizmodo.com/i-cut-google-out-of-my-life-it-screwed-up-everything-1830565500) <https://gizmodo.com/i-cut-google-out-of-my-life-it-screwed-up-everything-1830565500>
- [31\]](https://www.businessinsider.com/patreon-crowdfunding-platform-defends-itself-amid-boycott-2018-12) <https://www.businessinsider.com/patreon-crowdfunding-platform-defends-itself-amid-boycott-2018-12>
- [32\]](https://techcrunch.com/2019/01/29/facebook-project-atlas/) <https://techcrunch.com/2019/01/29/facebook-project-atlas/>
- [33\]](https://www.amazon.com/Freakonomics-Revised-Expanded-Economist-Everything/dp/0061234001) <https://www.amazon.com/Freakonomics-Revised-Expanded-Economist-Everything/dp/0061234001>
- [34\]](https://techcrunch.com/2019/01/30/googles-also-peddling-a-data-collector-through-apples-back-door/) <https://techcrunch.com/2019/01/30/googles-also-peddling-a-data-collector-through-apples-back-door/>
- [35\]](https://www.troyhunt.com/data-from-connected-cloudpets-teddy-bears-leaked-and-ransomed-exposing-kids-voice-messages/) <https://www.troyhunt.com/data-from-connected-cloudpets-teddy-bears-leaked-and-ransomed-exposing-kids-voice-messages/>
- [36\]](https://openconnectivity.org/blog/iot-security-design) <https://openconnectivity.org/blog/iot-security-design>
- [37\]](https://youtu.be/glFLpkCnSPU?t=237) <https://youtu.be/glFLpkCnSPU?t=237>
- [38\]](http://livingiot.cs.washington.edu/files/livingiot.pdf) <http://livingiot.cs.washington.edu/files/livingiot.pdf>
- [39\]](http://www.coolwearable.com/glucosentry-bracelet-diabetics/) <http://www.coolwearable.com/glucosentry-bracelet-diabetics/>
- [40\]](https://www.informationweek.com/healthcare/mobile-and-wireless/10-medical-device-wearables-to-improve-patients-lives/d/d-id/1323544) <https://www.informationweek.com/healthcare/mobile-and-wireless/10-medical-device-wearables-to-improve-patients-lives/d/d-id/1323544>
- [41\]](https://www.humanyze.com/) <https://www.humanyze.com/>
- [42\]](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1263992) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1263992
- [43\]](https://www.newscientist.com/article/mg21128191-600-specs-that-see-right-through-you/) <https://www.newscientist.com/article/mg21128191-600-specs-that-see-right-through-you/>
- [44\]](https://www.wired.com/story/eye-scanning-lie-detector-polygraph-forging-a-dystopian-future/) <https://www.wired.com/story/eye-scanning-lie-detector-polygraph-forging-a-dystopian-future/>
- [45\]](https://www.youtube.com/watch?v=JlO2X4Y96L8) <https://www.youtube.com/watch?v=JlO2X4Y96L8>
- [46\]](https://www.youtube.com/watch?v=Z4kqbKQrvYA) <https://www.youtube.com/watch?v=Z4kqbKQrvYA>
- [47\]](https://www.youtube.com/watch?v=75k8sqh5tfQ) <https://www.youtube.com/watch?v=75k8sqh5tfQ>
- [48\]](https://www.youtube.com/watch?v=iOucwX7Z1HU) <https://www.youtube.com/watch?v=iOucwX7Z1HU>
- [49\]](https://www.theguardian.com/technology/2018/may/24/amazon-alexa-recorded-conversation) <https://www.theguardian.com/technology/2018/may/24/amazon-alexa-recorded-conversation>
- [50\]](https://www.youtube.com/watch?v=j-KJxKHjb_w) https://www.youtube.com/watch?v=j-KJxKHjb_w
- [51\]](https://youtu.be/cUTMhmVh1qs?t=4654) <https://youtu.be/cUTMhmVh1qs?t=4654>
- [52\]](https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/) <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>
- [53\]](https://www.csoonline.com/article/3258748/security/the-mirai-botnet-explained-how-teen-scammers-and-cctv-cameras-almost-brought-down-the-internet.html) <https://www.csoonline.com/article/3258748/security/the-mirai-botnet-explained-how-teen-scammers-and-cctv-cameras-almost-brought-down-the-internet.html>

- [54\]](https://krebsonsecurity.com/2018/09/mirai-botnet-authors-avoid-jail-time/)https://krebsonsecurity.com/2018/09/mirai-botnet-authors-avoid-jail-time/
- [55\]](https://www.xkcd.com/936/)https://www.xkcd.com/936/
- [56\]](https://finance.yahoo.com/news/yahoo-data-breach-stolen-passwords-191113081.html)https://finance.yahoo.com/news/yahoo-data-breach-stolen-passwords-191113081.html
- [57\]](https://www.theregister.co.uk/2019/01/22/google_chrome_browser_ad_content_block_change/)https://www.theregister.co.uk/2019/01/22/google_chrome_browser_ad_content_block_change/
- [58\]](https://www.youtube.com/watch?v=YJg02ivYzSs)https://www.youtube.com/watch?v=YJg02ivYzSs
- [59\]](https://www.youtube.com/watch?v=9c6W4CCU9M4)https://www.youtube.com/watch?v=9c6W4CCU9M4
- [60\]](https://www.youtube.com/watch?v=-KmFSmkDyr8)https://www.youtube.com/watch?v=-KmFSmkDyr8
- [61\]](https://youtu.be/D7TB8b2t3QE?t=273)https://youtu.be/D7TB8b2t3QE?t=273
- [62\]](https://sites.google.com/site/glasscomms/glass-explorers)https://sites.google.com/site/glasscomms/glass-explorers
- [63\]](https://www.cnet.com/news/google-lens-google-glass/)https://www.cnet.com/news/google-lens-google-glass/
- [64\]](https://www.cnet.com/news/three-ways-google-maps-just-got-better/)https://www.cnet.com/news/three-ways-google-maps-just-got-better/
- [65\]](https://www.cnet.com/news/google-rolls-out-digital-wellbeing-tool-to-help-limit-screen-time/)https://www.cnet.com/news/google-rolls-out-digital-wellbeing-tool-to-help-limit-screen-time/
- [66\]](https://www.youtube.com/watch?v=Q4FoAr8i26g)https://www.youtube.com/watch?v=Q4FoAr8i26g
- [67\]](https://feelreal.com/)https://feelreal.com/
- [68\]](https://www.imdb.com/title/tt0104692/?ref_=fn_al_tt_4)https://www.imdb.com/title/tt0104692/?ref_=fn_al_tt_4
- [69\]](https://www.imdb.com/title/tt0133093/?ref_=fn_al_tt_1)https://www.imdb.com/title/tt0133093/?ref_=fn_al_tt_1
- [70\]](https://www.imdb.com/title/tt0181689/?ref_=fn_al_tt_1)https://www.imdb.com/title/tt0181689/?ref_=fn_al_tt_1